



**NATIONAL
OPTICAL
ASTRONOMY
OBSERVATORIES**

Preprint Series

NOAO Preprint No. 833

**MULTITAPER SPECTRAL ANALYSIS AND WAVELET
DENOISING APPLIED TO HELIOSEISMIC DATA**

R. W. KOMM
Y. GU
F. HILL
P. B. STARK
I. K. FODOR

Submitted to: The Astrophysical Journal

February 1999

Operated for the National Science Foundation by the Association of Universities for Research in Astronomy, Inc.

MULTITAPER SPECTRAL ANALYSIS AND WAVELET DENOISING APPLIED TO HELIOSEISMIC DATA

R.W. KOMM, Y. GU¹, and F. HILL

National Solar Observatory, National Optical Astronomy Observatories²,
950 N. Cherry Ave., Tucson, AZ 85726

and

P.B. STARK

Department of Statistics, Space Sciences Laboratory, and
Theoretical Astrophysics Center, UC Berkeley, CA 94720

I.K. FODOR

Department of Statistics, UC Berkeley, CA 94720

ABSTRACT

Estimates of solar normal mode frequencies from helioseismic observations can be improved by using Multitaper Spectral Analysis (MTSA) to estimate spectra from the time series, then using wavelet denoising of the log spectra. MTSA leads to a power spectrum estimate with reduced variance and better leakage properties than the conventional periodogram. Under the assumption of stationarity and mild regularity conditions, the log multitaper spectrum has a statistical distribution that is approximately Gaussian, so wavelet denoising is asymptotically an optimal method to reduce the noise in the estimated spectra. We find that a single m - ν spectrum benefits greatly from MTSA followed by wavelet denoising, and that wavelet denoising by itself can be used to improve m -averaged spectra.

We compare estimates using two different 5-taper estimates (Slepian and sine tapers) and the periodogram estimate, for GONG time series at selected angular degrees ℓ . We compare those three spectra with and without wavelet-denoising, both visually, and in terms of the mode parameters estimated from the pre-processed spectra using the GONG peak-fitting algorithm. The two multitaper estimates give equivalent results. The number of modes fitted well by the GONG algorithm is 20% to 60% larger (depending on ℓ and the temporal frequency) when applied to the multitaper estimates than when applied to the periodogram. The estimated mode parameters (frequency, amplitude and width) are comparable for the three power spectrum estimates, except

¹Present Address: NCR Corporation, 2651 Satellite Boulevard, Duluth, GA 30136

²Operated by the Association of Universities for Research in Astronomy, Inc. under cooperative agreement with the National Science Foundation.

for modes with very small mode widths (a few frequency bins), where the multitaper spectra broaden the modes compared with the periodogram.

At frequencies below 3 mHz, wavelet denoising of the log multitaper power spectra tends to increase the number of modes for which the GONG peak fitting algorithm converges well. Close to 3 mHz, where all modes are resolved, wavelet denoising makes little difference. At higher frequencies close to the acoustic cut-off frequency, where modes are blended into ridges, wavelet denoising the multitaper spectra reduces the number of good fits.

We tested the influence of the number of tapers used and found that narrow modes at low n values are broadened to the extent that they can no longer be fit if the number of tapers is too large. For helioseismic time series of this length and temporal resolution, the optimal number of tapers is less than 10.

Subject headings: Sun: oscillations — Techniques: time series analysis — Techniques: image processing

1. INTRODUCTION

The primary data products of helioseismology are the mode frequencies of acoustic oscillations, which are used to infer the structure and kinematics of the solar interior. With the excellent data available now from instruments such as GONG and SOHO-SOI/MDI, greater accuracy and reliability of the data processing is required to make substantial progress in the understanding of the solar interior (for example, the existence of a polar jet, as discussed by Howe et al. 1998). In addition, other mode parameters such as amplitude, asymmetry, and width are increasingly interesting. We address the step of converting the observed time series to frequency spectra, and study the potential benefits of modern time series analysis techniques. We apply Multitaper Spectral Analysis (MTSA) to the observed time series to derive power spectrum estimates, and then apply wavelet denoising to the log spectra to further improve the signal-to-noise ratio of the modes. MTSA has better bias and variance properties than the conventional periodogram and, under the assumption of stationarity and mild regularity, the log multitaper spectrum has approximately Gaussian statistics, so wavelet denoising is an asymptotically optimal method to reduce the noise level in the calculated spectra (cf. Walden, McCoy, & Percival 1995). Other wavelet-based methods to reduce noise in astronomical data exist, such as Murtagh, Starck, & Bijaoui (1995) and Fligge & Solanki (1997). We note that we use ‘standard’ multitaper techniques without special treatment of data gaps. This seems reasonable for GONG or SOHO-SOI/MDI data which have better than 85% duty ratio, and the gaps are short and more or less randomly distributed. Additional work (Fodor & Stark 1998) includes the gap structure and constructs optimal tapers for time series with known gaps. We have put together a ‘pipeline’ to calculate a multitaper spectral estimate from a given time series, apply wavelet denoising to the log spectra

and then derive mode parameters using the GONG peak-fitting algorithm of Anderson, Duvall, Jr., & Jefferies (1990). This pipeline was first applied to a set of simple artificial data to check for systematic errors and consistency, and then applied to observed time series of different lengths (daily, monthly, etc.).

We describe the method in detail and compare quantitatively three power spectrum estimates (5-taper estimates using Slepian or sine tapers, and the periodogram) using GONG month 16 times series for three different values of angular degree ℓ . We also compared each of the three spectrum estimates with the corresponding wavelet denoised spectrum. Multitapering helps quite generally, and wavelet denoising gives additional benefits at some frequencies. The new methods are not biased systematically compared with each other. The single best thing to improve mode fitting is to use a multitaper spectrum estimate. Wavelet denoising can further improve multitaper spectrum estimates for mode frequencies below about 3 mHz.

2. MULTITAPER SPECTRAL ANALYSIS

MTSA is an extension of single taper spectral analysis where the time series is multiplied/apodized with a single window function or data taper before calculating the power spectrum (Thomson 1982). Compared with the periodogram, a power spectrum estimate that uses a smooth window function, such as a Hanning window, can reduce spectral leakage (not to be confused with spatial leakage). The window functions that minimize leakage give less weight to the ends of the time series. The multitaper approach uses a variety of orthogonal tapers, some of which give more weight to the ends of the time series, trading off bias and variance. A multitaper estimate that uses well selected tapers can gain from the bias-variance tradeoff, giving an estimate that has small bias compared with a single taper estimate, but substantially lower variance. MTSA is described and motivated clearly and in detail by Percival & Walden (1993). The basic procedure is as follows: The time series $\{X_t\}$ is multiplied by each of K different data tapers, $\{h_{t,k}\}_{k=0}^{K-1}$. The periodogram of the K resulting series is computed, resulting in K estimates of the power spectrum, $\{\hat{S}_k^{(mt)}(f)\}_{k=0}^{K-1}$. The multitaper spectrum estimate is the average of these K power spectrum estimates:

$$\hat{S}^{(mt)}(f) \equiv \frac{1}{K} \sum_{k=0}^{K-1} \hat{S}_k^{(mt)}(f) \quad (1)$$

for

$$\hat{S}_k^{(mt)}(f) \equiv \Delta t \left| \sum_{t=0}^{N-1} h_{t,k} X_t \exp(-i2\pi f t \Delta t) \right|^2. \quad (2)$$

The tapers $\{h_{t,k}\}$ are normalized so that $\sum_{t=0}^{N-1} h_{t,k}^2 = 1$, and Δt is the sampling interval.

The tapers are chosen to be concentrated in the frequency domain, so that their broad band

bias is as small as possible. If the data tapers are pairwise orthogonal, i.e.,

$$\sum_{t=0}^{N-1} h_{t,j} h_{t,k} = 0 \quad \forall j \neq k, \quad (3)$$

then their corresponding power spectrum estimates, $\hat{S}_k^{(mt)}(f)$, are approximately uncorrelated (asymptotically uncorrelated for long time series for a wide variety of processes). The average of the K power spectrum estimates, $\hat{S}^{(mt)}(f)$, then has smaller variance than the individual power spectrum estimates by a factor that approaches $1/K$.³ Figure 1 gives an example of orthogonal tapers and shows the first five 4π Slepian tapers (4π discrete prolate spheroidal sequences) in the left column. The first taper resembles a conventional data taper such as the Hanning window: it gives more weight to the center of the time series than to its ends. Tapers for larger k give increasingly more weight to the ends of the time series. The right column shows the taper energy, which is the sum of the squared tapers, normalized by the number of tapers, K . It is evident that the portion of the time series that receives large weight increases as the number of tapers increases.

To ensure good leakage/bias properties, the tapers should be concentrated in the frequency domain. One way to define a family of ‘optimal’ tapers is as follows: (1) Specify a measure of concentration in the frequency domain. (2) Among all functions that vanish where there are no data, find the function for which the concentration is maximal. (3) Among all functions orthogonal to the first function and that vanish where there are no data, find the function for which the concentration is maximal. One repeats step (3) insisting that the function sought be orthogonal to all the previous functions found, until one has the desired number of tapers. For quadratic measures of spectral concentration, such as the ratio of the power in the band $[-W, W]$ to the total power, this can be cast as an eigenvalue problem (Slepian 1978, 1983). When observations are available continuously within an interval, and concentration is measured by the fraction of power in the band $[-W, W]$, the eigenfunctions are prolate spheroidal wavefunctions. For a discretized signal with observations available at every sample point within an interval, and concentration is measured by the fraction of power in the band $[-W, W]$, the eigenfunctions are called time-limited “discrete prolate spheroidal sequences” (dpss) or Slepian tapers. Figure 2 shows the first 11 eigenvalues, λ_k , of $n\pi$ Slepian tapers for $n = 3, 4$, and 5 with n being a measure of the resolution bandwidth, $2W = 2n/(N\Delta t)$, which increases with increasing n . The eigenvalues are the fraction of power in the given frequency interval. The closer λ_k is to 1, the less broad-band leakage. The first $(2n)$ tapers (with index $k = 0, \dots, n-1$) have eigenvalues larger than $\frac{1}{2}$ and the first $(2n-2)$ eigenvalues are extremely close to 1: the tapers are nearly perfectly concentrated to the band $[-W, W]$. The last $N-2n$ tapers have most of their power outside the band $[-W, W]$ and thus have poor broad-band bias. The number of tapers with good leakage properties increases with n and with the concentration bandwidth $2W$. Typically, n is chosen to be 2, 3, or 4 to limit

³However, this does not lead to a proportional reduction in the formal errors of fitted mode parameters, see Section 5.2.

bias caused by the width of the central lobe of the power spectrum of the taper. The trade-off depends on the length of the time series and the frequency resolution required for subsequent data processing, among other things.

Riedel & Sidorenko (1995) introduced a different measure of spectral concentration for tapers, based on an asymptotic expression for the local bias. Their measure again leads to a quadratic optimization problem that can be cast as an eigenvalue problem. They showed that the optimal tapers for that definition of spectral concentration were approximated quite well by sine tapers, which are extremely easy to compute, and are orthogonal (for data without gaps). The k -th sine taper is

$$v_i^k = \sqrt{\frac{2}{N+1}} \sin\left(\frac{\pi k i}{N+1}\right) \quad (4)$$

with $i = 0, \dots, N-1$. The multiplicative constant makes the tapers orthonormal.

For stationary Gaussian processes, the distribution of a singly tapered spectrum estimate at a given frequency is that of a constant times a χ_2^2 random variable. In contrast, if K orthogonal tapers are used, the distribution is approximately (asymptotically in the length of the time series) that of a constant times a χ_{2K}^2 random variable. The logarithm of the multitaper spectrum estimate has a more nearly Gaussian distribution than does the logarithm of a singly tapered estimate. Indeed, Walden, McCoy, & Percival (1995) find that if five or more tapers are used, the Gaussian approximation is good, although that result clearly depends on details of the underlying spectrum and the length of the time series. Wavelet shrinkage denoising is asymptotically an optimal method to enhance the signal-to-noise ratio of the estimate of a function observed with additive Gaussian noise (Donoho et al. 1993); Walden, McCoy, & Percival (1995) demonstrate its application to multitaper spectrum estimates. Wavelet shrinkage denoising is a nonlinear procedure that does not tend to smear fine-scale features in the data, as linear smoothing does.

3. WAVELET DENOISING

We use the discrete wavelet transform (cf. Donoho & Johnstone 1994 and references therein) with orthogonal wavelet bases; examples of orthogonal wavelet bases are the original Daubechies wavelets and the Coiflets and Symmlets, cf. Daubechies (1993). Wavelet denoising works for many classes of signals because, as reasonable models of real-world phenomena, they usually have extremely sparse wavelet decompositions: relatively few large coefficients yield an excellent approximation. On the other hand, Gaussian white noise remains Gaussian with the same rms-amplitude after a wavelet transform. The basic strategy in wavelet denoising is to define a thresholding or “shrinkage” scheme that retains the large wavelet coefficients, which are predominantly signal, and rejects the small coefficients, which are predominantly noise. For a formal argument see, for example, Donoho (1993).

The discrete wavelet transform maps a signal of $N = 2^{J+1}$ discrete data points into the same

number of wavelet coefficients, w_{jk} , with $(J + 1)$ scales or dyads, indexed by $j = 0, \dots, J$, and (2^j) coefficients per dyad, indexed by $k = 0, \dots, 2^j - 1$. Figure 3 shows different thresholding schemes such as hard (keep or kill) and soft (shrink or kill). Hard thresholding keeps all wavelet coefficients above a certain threshold, and sets all smaller coefficients to zero, while soft thresholding reduces even the coefficients above the threshold by the amount of the threshold. The threshold, θ given by Donoho & Johnstone (1994) is

$$\theta \equiv \sqrt{2 \log(N)} \sigma_j, \quad (5)$$

where N is the signal length and σ_j is the noise level in the observations at scale j . In level-independent shrinkage, σ_j is estimated by the median absolute deviation (MAD) of the wavelet coefficients at the finest scale (J), normalized by 0.6745 to correspond to the standard deviation of a Gaussian distribution:

$$\hat{\sigma}_j \equiv \text{median}(|w_{Jk} - \text{median}(w_{Jk})|)/0.6745. \quad (6)$$

In level-dependent shrinkage, the standard deviation of the noise level at scale j is estimated by the median absolute deviation of the coefficients at scale j ; because the coefficients at broad scales tend to have nontrivial components of the signal, this can over-estimate the noise level at those scales, resulting in too much shrinkage, attenuating the real signal. In tests on artificial and real helioseismic data, we found that level-dependent shrinkage distorted the modes unacceptably. This agrees with Walden, McCoy, & Percival (1995) who found that with the exception of white noise, scale-independent thresholding leads to better results than scale-dependent thresholding. Therefore, we use level-independent shrinkage throughout this work.

Hard thresholding recovers the signal well in mean squared error, but tends not to suppress some noise spikes (spikes do not make a large contribution to mean squared error). Soft thresholding leaves fewer noise spikes, but tends to attenuate the signal because even the largest wavelet coefficients are shrunk towards zero. Gao & Bruce (1995) introduced semisoft thresholding (kill, shrink, or keep) as a compromise between hard and soft thresholding. The two thresholds defining three ranges bracket θ as defined in Equation (5) and are given as a function of N in Table 1 of their paper.

In tests using artificial and real data, we found that a modified semisoft or “semihard” thresholding scheme worked best for helioseismic data. The threshold function is shown in Figure 3. The lower threshold is θ as defined in Equation (5), and the upper semisoft threshold is that of Gao & Bruce (1995). The semihard thresholding scheme reduces the visual roughness of the estimate more than the semisoft scheme without distorting the signal structure. Hard thresholding also preserves the mode structure, but gives estimates that are visually too rough, while soft thresholding gives the smoothest estimates, but broadens the modes unacceptably. We use the semihard threshold below.

Our ultimate wavelet shrinkage procedure is the following: we wavelet-transform the log power spectrum, scale the wavelet coefficients by the estimated deviation, $\hat{\sigma}_J$, of the coefficients

at the finest scale, apply the “semi-hard” threshold to the coefficients, rescale the coefficients, and inverse wavelet-transform them to obtain the “denoised” estimate of the log power spectrum.

This denoising scheme can introduce a small systematic frequency shift due to the lack of translation invariance of the wavelet basis. To eliminate this effect, we used translation-invariant wavelet denoising (Coifman & Donoho 1995) which efficiently shifts the signal over all positions and averages the denoised shifted signals (after shifting them back). This procedure also reduces the influence of the specific wavelet basis chosen. We tried Haar, Daubechies, Coiflet, and Symmlet bases. We found that for our data, Symmlets of order 8 produced estimates that were visually preferable, and took the least time to compute.

The denoising procedure can be restricted to a range of scales from the smallest to an upper limit. It might not be necessary to denoise the largest scales, because they tend to contain primarily large-scale trends in the signal and not noise. A value of $j_{\min} = 4$ or $j_{\min} = \text{int}(\sqrt{2 \log(N)})$ is a good choice for signals of length $N = 2048$ or less, while a larger value of $j_{\min} \geq 7$ is more appropriate for long helioseismic data sets, where N is typically about $N = 2^{16}$ or larger.

4. IMPLEMENTATION

To calculate multitaper spectrum estimates from helioseismic time series, we used 3π Slepian tapers to ensure that the widths of p-modes are not broadened by the resolution bandwidth of the tapers. We averaged over $K = 5$ singly tapered spectrum estimates. We also calculated multitaper spectrum estimates using the first five sine tapers. For MTSA with Slepian tapers, we used subroutines written in C by Lees & Park (1995). To create the multitaper spectrum estimates, we averaged the singly tapered spectrum estimates weighted according to their respective eigenvalues. To wavelet denoise the log-spectra, we used the WaveLab package by Buckheit et al. (1995), translated to IDL by Graps (1995). We added translation-invariant denoising, which is part of the current version of WaveLab, to the IDL package and implemented semisoft and semihard thresholding, which were not in the IDL or WaveLab package. To derive mode parameters from the resulting spectra, we used the GONG peak-fitting algorithm developed by Anderson, Duvall, Jr., & Jefferies (1990).

The peak-fitting algorithm uses a maximum-likelihood method assuming χ^2 statistics with two degrees of freedom, while the multitaper estimate asymptotically drives the analyzed spectra toward Gaussian statistics (with k tapers leading to χ^2 statistics with $2k$ degrees of freedom). General minimization algorithms assume either Gaussian or χ^2 statistics. Anderson, Duvall, Jr., & Jefferies (1990) mention that the changes required to fit an average spectrum of k spectra is to multiply the likelihood function by k and divide the error by \sqrt{k} . They compared the results of fitting 1,000 artificial realizations (of a single spectrum) using their maximum likelihood method and using a least-squares approach. They found that both techniques yield the same mode frequency, but that only the maximum likelihood technique yields reliable estimates of mode

width, amplitude and background in this case. Thus, we proceed using the GONG peak-fitting algorithm and compare mode parameters derived from the different spectral estimates in order to check for systematic differences.

We applied the procedures to artificial data to check for systematic frequency shifts. To test multitaper spectral analysis, we created a set of time series with a known limit spectrum, following Schou & Brown (1992) to model individual modes as stochastically excited damped oscillators. We assume that the mode lifetime is small compared to the length of the whole time series but large compared to the time cadence of observation. The mode is randomly excited many times over the whole time interval with the time interval between two ‘kicks’ small compared to the mode lifetime. To simplify the modeling, we excite the mode only at temporal grid points. Two or more kicks at the same time grid point are considered to be one kick. We compensate for overlapping kicks by adjusting the sample size of uniformly distributed random numbers. The amplitude of each kick is scaled to ensure that the average power stays constant. In this way, we generate noise-free time series of a single mode. Multi-mode time series are generated by summing several single-mode time series. This ignores coupling in the excitation process or by nonlinearity. We calculated 1,000 realizations of the $n = 10$, $\ell = 50$ mode (with a length of 36 days, one ‘GONG month’). For each time series, we calculated a periodogram, a 5-taper 3π Slepian spectrum estimate, and a 5-sine taper spectrum estimate. We then used the GONG peak-fitting procedure to estimate mode parameters from the various spectrum estimates. The average mode frequency differs by 2 nHz with a standard deviation of 226 nHz (periodogram), 1 ± 233 nHz (Slepian multitaper), and 1 ± 245 nHz (sine multitaper) from the input mode frequency ($\nu = 3045.5339 \mu\text{Hz}$). The standard deviation is slightly smaller than one frequency bin. The rms differences between periodogram and multitaper spectra are 44 nHz (Slepian) and 51 nHz (sine). The frequency estimates determined from the periodogram and multitaper spectra do not differ systematically in a statistically significant way.

To evaluate wavelet denoising, we created artificial spectra that contained a single mode plus realization noise for four different signal-to-noise ratios ($\text{S/N} = 50, 20, 10$, and 5). We modeled the modes as Lorentzian profiles using a subroutine written by E. Anderson. For each signal-to-noise ratio, we created 1,000 artificial spectra with a spectral pixel size of $0.2 \mu\text{Hz}$ and a mode frequency of $\nu = 3045.9 \mu\text{Hz}$, which is close to but slightly different from the $n = 10$, $\ell = 50$ mode. We applied single wavelet denoising and translation-invariant (TI) denoising to the spectra, using semihard thresholding. For large signal-to-noise ratio ($\text{S/N}=50$), we find that the average mode frequency differs by -6 nHz with a standard deviation of 212 nHz (original), -5 ± 214 nHz (denoised), and -8 ± 212 nHz (TI denoised) from the input mode frequency. The rms difference between the denoised and original spectra is 13 nHz, compared with only 3 nHz for TI denoised and original spectra. For small signal-to-noise ratio ($\text{S/N}=10$), the average difference is -65 ± 306 nHz for the simple denoising scheme, while it is -10 ± 303 nHz for the TI-denoised spectra and -12 ± 283 nHz for the original spectra. The rms differences between the denoised and original spectra increase to 107 nHz (denoised) and to 46 nHz (TI denoised). Simple wavelet

denoising can introduce a small, systematic shift in frequency, while TI denoising does not show this systematic shift. Therefore, we use TI wavelet denoising henceforth in this study.

5. RESULTS

5.1. Helioseismic Data: GONG month 16

We applied MTSA and wavelet denoising to different helioseismic data sets obtained from South Pole, SOI and GONG. For quantitative comparisons, we use a 36-day GONG time series for angular degrees $\ell = 30, 65$, and 100 . A merged GONG month time series is created by combining 36 network days; GONG month 16, used here, begins 28 Oct. 1996 and ends 2 Dec. 1996. GONG has a one-minute cadence, so the network month time series contains 51,840 data points. For month 16, the data fill factor is 0.94; short gaps were filled using an autoregressive filter, gaps larger than 2 minutes are set to zero. The data set for each ℓ contains the complex time series for each spherical harmonic, m , from 0 to ℓ , and in addition the gap structure (before and after gap-filling) as a function of time.

Figure 4 shows three power spectrum estimates of the GONG month 16 velocity time series (28 Oct 1996 – 2 Dec 1996, 36 days) of $\ell = 100$, $m = 0$ in the left column and the corresponding wavelet denoised spectra in the right column. For simplicity, we show only a small part of the spectra; the complete spectra extend over the bins from 0 to $8333 \mu\text{Hz}$. For the MTSA spectrum estimates, we used five 3π Slepian tapers and five sine tapers. The Slepian multitaper spectrum is clearly better than the periodogram. It has much smaller variance, and the modes are better resolved. The sine multitaper spectrum estimate is quite similar to the Slepian estimate except for some background details.

The right column in Figure 4 shows that wavelet denoising cleans the periodogram, but it reduces the mode amplitude drastically. The ‘click’ at about $3000 \mu\text{Hz}$ in the upper right panel is a typical artifact of noise in one wavelet coefficient exceeding the threshold, and thus remaining unattenuated in the denoised reconstruction. Wavelet denoising works well for the two multitaper spectra, leaving the principal mode and leaks clearly distinguishable, and smoothing the mode structure without broadening it.

5.2. Quantitative Comparison

We compared the number of primary modes the GONG peak-fitting algorithm fits successfully for the different spectrum estimates. The peak-fitting algorithm has two error flags related to the quality of the fit to a mode. One is based on heuristic assumptions about the modes; the other indicates numerical difficulties (cf. Hill et al. 1998). The numerical flag indicates how well the minimization of the likelihood function converges and distinguishes between failure to converge,

convergence with some difficulty, and strict convergence is necessary but not sufficient for a good fit. The top row of Figure 5 shows an example where the fit to a periodogram fails to converge (left panel), while the fit to the corresponding denoised multitaper spectrum converges (right panel). The heuristic flag includes, for example, a test to ensure that the fit has not locked onto the first guess; that the fitted width is within a factor of two of the first guess width; etc. The bottom row of Figure 5 shows an example where the fit to a periodogram is rejected by the heuristic flag (left panel) because the fitted frequency error is larger than half the first guess width. The corresponding denoised multitaper spectrum leads to a good fit (right panel). We use only mode fits that are good according to both flags.

The three panels in Figure 6 shows histograms of the number of good fits (solid line) as a function of frequency for the periodogram, the sine multitaper, and corresponding denoised multitaper spectrum estimates for angular degree $\ell = 30$. The corresponding histograms of the Slepian multitaper spectra are very similar and are not shown here. Each bin corresponds to all modes of a single radial order n averaged over all angular orders m values. The radial orders of the modes in the data range from $n = 4$ to $n = 25$. The total number of fits, good and bad, is shown by the dotted line. In the frequency range from about 2000 to 4000 μ Hz, all $(2\ell + 1)$ modes are present in the data. For the periodogram in the top panel, the number of good fits is only 43% of the total number of modes. This fraction increases to 80% for the multitaper spectrum estimate in the middle panel. With MTSA, the number of good fits increases for all frequency bins compared to the periodogram results (included as dashed line). The bottom panel shows the wavelet-denoised multitaper spectrum estimate (solid line) and the multitaper spectrum estimate (dashed line) for comparison. Wavelet denoising further increases the number of good fits at frequencies below about 2000 μ Hz, while it tends to reduce the number of good fits at frequencies above about 4000 μ Hz, where modes blend into ridges. In the frequency range of well-resolved modes around 3000 μ Hz, wavelet denoising increases the number of good fits at some n values and reduces it at others, leading to a small net gain.

Figures 7 and 8 show the same for $\ell = 65$ and $\ell = 100$. As in Figure 6, each bin contains all modes of a single n value ranging from $n = 1$ to $n = 19$ for $\ell = 65$, and from $n = 0$ to $n = 15$ for $\ell = 100$. Figure 7 shows that for $\ell = 65$, multitapering greatly improves the number of good fits at all frequencies compared to the periodogram, but especially between about 2000 and 4000 μ Hz, where almost all possible $(2\ell + 1)$ modes can be fitted well to the multitaper spectrum estimates. Wavelet denoising improves the fits at low n values (below about 2000 μ Hz). The fit of the $n = 1$ mode is good only for the denoised spectrum estimate. Denoising makes no difference in the range of well-resolved modes, where multitapering leads already to good fits at all modes, while at high frequencies denoising reduces the number of good fits. In Figure 8, the $\ell = 100$ spectra spectrum estimates show a similar behavior.

Table 1 shows the total number of fitted modes and the number of good fits for angular degrees $\ell = 30, 65$, and 100 (cf. Figures 6 to 8). The numbers of good fits are separated into three frequency ranges: (1) $\nu \leq 2.5\text{mHz}$, low signal-to-noise modes; (2) $2.5\text{mHz} < \nu < 3.5\text{mHz}$,

well-resolved modes; and (3) $\nu \geq 3.5\text{mHz}$, blended modes. Not surprisingly, the total number of primary modes is about the same for all three estimates (cf. dotted line in Figures 6 to 8). However, the number of good fits increases by about 60% when a multitaper spectrum estimate is fitted instead of the periodogram. As a function of frequency, the smallest increase is in the frequency range (2), where the modes are well-resolved ($\approx 12\%$ at $\ell = 65$), since there the number of good fits does not depend strongly on the method used. The largest increase is in the low frequency range (1) ($\approx 100\%$ at $\ell = 65$), where the modes have a low signal-to-noise ratio, while the increase is average in the high frequency range (3) ($\approx 30\%$ at $\ell = 65$), where the modes are blended into ridges.

Table 2 is the same as Table 1, but for the wavelet-denoised spectrum estimates for angular degrees $\ell = 30, 65$, and 100 . As expected from Figure 4, denoising oversmooths the periodogram and reduces the number of good fits substantially, for example, to 65% of the value in Table 1 for $\ell = 65$. For the multitaper spectrum estimates, the number of good fits increases by 5% at low and mid frequencies compared to the values in Table 1, while at high frequencies, the number of good fits increases by 4% for $\ell = 30$ and decreases by 25% at $\ell = 100$.

Figure 9 compares three mode parameter estimates (frequency, full-width at half maximum, and amplitude) of all good fits common to the three power spectrum estimates for $\ell = 65$, between 2.5 mHz and 3.5 mHz . At this ℓ value and frequency range, almost every mode is fitted in all three spectral estimates (cf. Figure 7). The rows show estimated mode frequency (ν), width (Γ), and amplitude (A) from top to bottom. The left column presents scatter plots of periodogram parameter estimates (x-axis) versus Slepian spectrum parameter estimates (y-axis), the middle column shows periodogram (x-axis) versus sine multitaper spectrum estimates (y-axis), and the right shows Slepian (x-axis) versus sine (y-axis) spectrum estimates. The initial guess was subtracted from the frequencies. The two background parameters, background amplitude and slope, are closely linked to the mode amplitude, and show a behavior similar to the mode amplitude. The three parameters show a positive correlation close to one between each pair of estimates. The scatter is the smallest between the two multitaper estimates in the right column.

To quantify this correlation, we calculated a linear regression between the parameters of any two spectrum estimates for each of the three mode parameters taking into account the errors in both data sets. The regression is included in Figure 9 as solid line. The regression parameters are tabulated in Table 3. For the mode frequency, the regression slope is very close to one and the intercept is below 10 nHz , determined from all 1285 good fits common to the three spectra. As a function of frequency, this is also true for modes below 3.5mHz . For modes above 3.5 mHz , the slope is about 0.84 and the intercept is as large as 32 nHz between multitaper spectra and periodogram.

For mode width, the slope of the regression line for the multitaper spectra is close to one, independent of frequency, and between the periodogram and any of the two multitaper spectra in the range of well-resolved modes ($2.5\text{ mHz} < \nu < 3.5\text{ mHz}$). The intercept is relatively small

in these cases and consistent with zero. For the low- and high-frequency ranges, regressing multitaper spectra against the periodogram leads to a slope of about 0.8 and a positive intercept. At frequencies above 3.5 mHz, the modes are rather broad, and with increasing width, the scatter increases to the point that the width estimated from the periodogram is no longer linearly associated with the width estimated from a multitaper spectrum, making a linear regression meaningless. If the regression is limited to modes where the mode width is smaller than half the distance to the nearest ℓ leak ($\frac{d\nu}{d\ell} > 2\Gamma$), the slope increases to 0.88 between estimates from periodogram and from the two multitaper spectrum estimates; the intercept is reduced by a factor of three. In addition, the corresponding regression slope between the frequencies increases from 0.84 to 0.99 and the intercept is reduced by 10 nHz. The widths of the modes below 2.5 mHz are less than 1.60 μ Hz (about six frequency bins), which means that they are comparable in size to the width of the central lobe of the combined tapers (cf. Section 2). Thus, these narrow modes are slightly broadened by the multitapering.

For the mode amplitude, the regression does not show a frequency dependence. The slope is close to one between every pair of power spectrum estimates, and the intercept is zero compared to the mode amplitudes. The background amplitude and the background slope, not shown here, give regression results similar to those of the mode amplitude.

Figure 10 shows a comparison of formal parameter fit errors determined from the inverse of the Hessian matrix at the fitted parameter values ($\delta\nu$: frequency, $\delta\Gamma$: width, δA : amplitude) for the well-resolved modes. It is known that the formal error is an unreliable estimate of the true uncertainty and reproducibility of the parameter estimates. As in Figure 9, the left column shows periodogram versus Slepian spectrum, the middle column shows periodogram versus sine spectrum, and the right shows Slepian versus sine spectrum. The frequency and width errors are in μ Hz. The average frequency error is about 0.2 μ Hz for modes below 3.5 mHz, comparable to but slightly smaller than a single frequency bin in the spectra, and 0.9 μ Hz for modes above 3.5 mHz. As with the mode parameters, the mode errors are positively correlated, with a slope close to one and a small, nearly zero, intercept. The different spectral estimates lead essentially to the same formal fit errors.

The peak-fitting algorithm converges in fewer iterations for the multitaper spectrum estimates than for the periodogram. The average number of iterations is 75.1 ± 20.8 for the periodogram, 57.0 ± 11.9 for the Slepian and 56.7 ± 11.8 for the sine multitaper spectrum. This is a reduction of 24%.

We calculated the difference between parameters from any two power spectrum estimates and scaled them by the formal error estimate provided by the peak-fitting algorithm, which is another way to check for systematic offsets. Figure 11 shows histograms of fit parameter differences scaled by the formal fit errors (ν : frequency, Γ : width, A : amplitude) for well-resolved modes between 2.5 and 3.5 mHz. The left column shows Slepian spectrum estimate minus periodogram scaled by periodogram error, the middle column shows sine spectrum estimate minus periodogram scaled

by periodogram error, and the right one shows sine minus Slepian spectrum estimates scaled by Slepian error. The distributions are all well within one formal error bar centered around zero, and the multitaper spectrum estimates (right column) lead to the narrowest distributions, as in Figure 9. Offsets are small compared to the fit error and are on the order of percent ($2\% \pm 1\%$) for all three mode parameters. The same is true for modes at lower and higher frequencies except for the mode width between periodogram and multitaper spectrum (cf. Table 3), where the offset is about 20% of the error.

Wavelet denoised spectrum estimates show the same general behavior. A scatterplot of estimated mode parameters of all good fits common to the periodogram and either of the two denoised multitaper spectra looks very similar to those shown in Figure 9. We performed a regression of the parameters estimated for every pair of spectrum estimates. In Table 4, the regression slopes and intercepts show the same behavior as in Table 3, with regression slopes close to one and intercepts close to zero. We repeated this analysis for $\ell = 30$ and $\ell = 100$ and found the same result as for $\ell = 65$.

5.3. Number of Tapers

As discussed in Section 2, there is a trade-off between bias and variance in choosing the number of tapers to use in MTSA. If the number of tapers is too large, detail in the spectrum estimates is lost, while if it is too small, the variance remains unnecessarily large. Resolving modes of a given width sets an upper limit for the number of tapers (see mode width for $\nu \leq 2.5$ mHz in Table 3). To estimate this limit, we calculated multitaper spectrum estimates using up to 50 tapers for $\ell = 65$. Figure 12 shows that the modes broaden with increasing number of tapers. As a consequence, the number of good fits decreases when more than 10 tapers are used, to about 28% for $K = 50$. The decrease is frequency-dependent; modes with small widths are most sensitive to the broadening. When the number of tapers increases, modes at increasingly higher frequencies can no longer be fit. For example, all modes at the lowest n value present in the data ($n = 1$) disappeared when the number of tapers increased from 5 to 10 and for 50 tapers only modes of $n = 10$ and higher can be fit. This test suggests that for helioseismic time series of this length, frequency resolution, and gap structure, the optimal number of tapers is below 10.

To study in more detail the influence of the number of tapers on the quality of estimated mode parameters, we repeated the analysis using 4 to 10 tapers and calculated the number of good fits as a function of n averaged over all m values. For $n = 3$ to 13, the number of tapers had a negligible effect on the number of good fits: all spectra led on average to $98\% \pm 2\%$ good fits of the $(2\ell + 1)$ possible modes. For $n = 14$ to 18, the number of good fits varies with the number of tapers, and generally increases with the number of tapers. However, this is the frequency range where modes are blended into ridges ($\frac{d\nu}{d\ell} < 2\Gamma$) and the current GONG peakfinding algorithm should not be applied. For $n = 1$ and 2, the number of good fits decreases with increasing number of tapers; using 9 and 10 tapers leads to substantially smaller numbers of good fits. This suggests

that the number of good fits is nearly constant for $k = 4$ to 7. Using a small number of tapers leaves the computations inexpensive and avoids excessive broadening of modes present for small n .

6. CONCLUSION

(1) Multitapering and wavelet denoising allow the GONG peak-fitting algorithm to fit more mode parameters successfully, as defined by the error flags in the GONG fitting procedure (cf. Tables 1 and 2). The improvement depends on the angular degree ℓ , the temporal frequency, and details of the time series. For the time series used in this work, multitapering increases the number of modes fitted well by 20% to 60% at frequencies below 3.5 mHz. Wavelet denoising adds about 5% more. We have also analyzed multitaper spectra covering $\ell = 0$ to 150 of a 3-month GONG time series and found that the improvement is 10% on average over all ℓ values and frequencies from 1 to 5 mHz with the largest improvement for $\ell \leq 70$ (cf. Komm et al. 1998).

(2) Multitapering and wavelet denoising do not lead to smaller formal error bars as computed by the peak-fitting algorithm (cf. Figure 10). Both methods reduce the variance or noise in the spectrum, making it easier to fit the modes in the spectrum (the peak-fitting algorithm takes fewer iterations). It is not known whether the effective uncertainty in the mode parameters is smaller for multitaper spectrum estimates and wavelet-denoised spectrum estimates, because the formal error bars do not measure the true reproducibility and uncertainty very well.

(3) Multitapering and translation-invariant wavelet denoising do not introduce any obvious systematic changes in the estimates of mode parameters (cf. Figures 9 and 11). Narrow modes at low frequencies are broadened by multitapering by about 20%, but these modes are not well-resolved in the first place. The modes are not broadened by wavelet denoising, as they would be by a linear filter.

(4) The benefit of multitapering occurs for a small number of tapers. The number of tapers has to be small in order not to broaden narrow modes at low frequencies excessively. For well-resolved modes at higher frequencies, the number of good fits does not depend strongly on the number of tapers (see Section 5.3). The ‘optimal’ number of tapers depends on the length of the time series. For example, for a 3-month GONG time series, we found that using 7 tapers leads to the largest number of good mode fits, but increases the mode widths only slightly (by $2.9\% \pm 1.8\%$ on average for $\ell = 65$, $n = 3\text{--}13$).

(5) Considering both techniques discussed here, the largest improvement at the smallest computational cost results from using multitaper spectrum estimates with sine tapers. Both multitapering and wavelet denoising improve mode fitting, but the larger improvement results from multitapering, which requires less computational effort. Because the difference between spectrum estimates using Slepian tapers and sine tapers is negligible for these data, it is adequate to use the more easily and inexpensively computed sine tapers. This is also the case for generalized sine tapers that take the gap structure into account (Fodor & Stark 1998).

We recommend using multitaper spectrum estimates at all frequencies and applying wavelet denoising to multitaper spectra at frequencies below 3.0 mHz.

This work was supported by NASA/SOHO/SOI/Stanford and by NSF Grant AST-9504410, NASA Grant NAG 5-5035, and NASA Grant NAG 5-3941. This work utilizes data obtained by the Global Oscillation Network Group (GONG) project, managed by the National Solar Observatory, a Division of the National Optical Astronomy Observatories, which is operated by AURA, Inc. under a cooperative agreement with the National Science Foundation. The data were acquired by instruments operated by the Big Bear Solar Observatory, High Altitude Observatory, Learmonth Solar Observatory, Udaipur Solar Observatory, Instituto de Astrofísico de Canarias, and Cerro Tololo Interamerican Observatory. To calculate multitaper spectra, we used subroutines written in C by Lees & Park (1995) and for the wavelet analysis we used WaveLab by Buckheit et al. (1995) translated to IDL by Graps (1995).

REFERENCES

Anderson, E.R., Duvall, Jr., T.L., and Jefferies, S.M. 1990, *ApJ*, 364, 699

Buckheit, J., Chen, S., Donoho, D., Johnstone, I., & Scargle, J. 1995, Technical Report, Dept. of Statistics, Stanford University

Coifman, R.R. & Donoho, D.L. 1995, in *Wavelets and Statistics*, Antoniadis, A. & Oppenheim, G. (eds.), 125, Springer Verlag, New York

Daubechies, I. 1993, *SIAM J. Math. Annal.*, 24, 499

Donoho, D. 1993 Stanford University Tech. Rep. 416

Donoho, D., Johnstone, I., Kerkyacharian, G. & Picard, D. 1993, *J. Roy. Statist. Soc. Series B*, 2, 301

Donoho, D. & Johnstone, I. 1994, *Biometrika*, 81, 425

Fligge, M. & Solanki, S.K. 1997, *A&AS*, 124, 579

Fodor, I. & Stark, P. 1998, *SOHO 6/GONG 98*, Ed. Korzennik, S., ESA Special Publication SP-418

Gao, H.-Y. & Bruce, A.G. 1995, *StatSci Research Report No. 39*

Graps, A. 1995, *The IDL Wavelet Workbench*, Research Systems, Inc, Boulder, CO

Hill, F., Anderson, E., Howe, R., Jefferies, S.M., Komm, R.W., & Toner, C.G. 1998, *Dynamics of the Interior of the Sun and Sun-like Stars* (Eds. S.G. Korzennik & A. Wilson), ESA SP-418, ESA Publications Division, Noordwijk, The Netherlands, 1998

Howe, R., Antia H., Basu S., Christensen-Dalsgaard J., Korzennik S.G., Schou J., and Thompson M.J. 1998, Dynamics of the Interior of the Sun and Sun-like Stars (Eds. S.G. Korzennik & A. Wilson), ESA SP-418, ESA Publications Division, Noordwijk, The Netherlands, 1998

Komm, R.W., Anderson, E., Hill, F., Howe, R., Fodor, I. & Stark, P. 1997, Dynamics of the Interior of the Sun and Sun-like Stars (Eds. S.G. Korzennik & A. Wilson), ESA SP-418, ESA Publications Division, Noordwijk, The Netherlands, 1998

Lees, J.M. & Park, J. 1995, Computer and Geology, 21, 199

Murtagh, F., Starck, J.-L., & Bijaoui, A. 1995, A&AS, 112, 179

Percival, D.B. & Walden, A.T., 1993, Spectral Analysis for Physical Applications, multitaper and conventional univariate techniques, Cambridge University Press, Cambridge

Riedel, K.S. & Sidorenko, A. 1995, IEEE Trans. Signal Process., 43, 188

Schou, J. & Brown, T.M. 1992, in Dissertation by J. Schou, p.99

Slepian, D. 1978, Bell Syst. Tech. J., 57, 1371

Slepian, D. 1983, SIAM Review, 25, 379

Thomson, D.J. 1982, Proceedings of the IEEE, 70(9), 1055

Walden, A.T., McCoy, E.J., & Percival, D.B. 1995 Technical Report, Dept. of Mathematics, Imperial College of Science Technology and Medicine, 28

Fig. 1.— The first five 4π Slepian tapers (discrete prolate spheroidal sequence data tapers) and the taper energy as a function of t for different number of tapers ($K = 1$ to 5).

Fig. 2.— The first 11 eigenvalues, λ_k , of $n\pi$ Slepian tapers for $n = 3, 4$, and 5 with n being a measure of the resolution bandwidth. The resolution bandwidth $2W$ is defined as $W = n/(N\Delta t)$ with N being the signal length and Δt the temporal resolution.

Fig. 3.— Thresholded wavelet coefficients as a function of wavelet coefficients for four different thresholding schemes shown for a signal of length $N = 2048$.

Fig. 4.— Three power spectrum estimates of the GONG month 16 velocity time series (28 Oct 1996 – 2 Dec 1996, 36 days) of $\ell = 100$, $m = 0$ and the corresponding wavelet denoised spectra.

Fig. 5.— Two examples where the fit to a periodogram fails (left panels) and the corresponding denoised multitaper spectrum leads to a good fit (right panels). The thick solid lines indicate the fits of the multitaper spectra and the dotted lines indicate the central frequency of the fitted mode. Top row: Failure to converge (numerical flag), $\ell = 65$, $m = -64$, $n = 9$. Bottom row: Fitted frequency error is larger than half the first guess width (heuristic flag), $\ell = 65$, $m = -61$, $n = 4$.

Fig. 6.— Histograms of the number of good fits (solid line) as a function of frequency for the periodogram, the sine multitaper spectrum, and the corresponding denoised multitaper spectrum of $\ell = 30$. Each bin contains the modes of a single n value summed over all m . The dotted line represents the total number of fits including the bad ones, and the dashed line represents the good fits from the panel above.

Fig. 7.— As Figure 6, but for $\ell = 65$.

Fig. 8.— As Figure 6, but for $\ell = 100$.

Fig. 9.— A comparison of fitted mode parameters (ν : frequency, Γ : width, and A : amplitude). The left column shows periodogram versus Slepian spectrum, the middle column shows periodogram versus sine spectrum, and the right shows Slepian versus sine spectrum. For ν , the initial guess value was subtracted.

Fig. 10.— A comparison of fit parameter errors ($\delta\nu$: frequency, $\delta\Gamma$: width, δA : amplitude). The left column shows periodogram versus Slepian spectrum, the middle column shows periodogram versus sine spectrum, and the right shows Slepian versus sine spectrum.

Fig. 11.— Histograms of fit parameter differences scaled by the fit errors (ν : frequency, Γ : width, A : amplitude). The left column shows Slepian minus periodogram scaled by periodogram error, the middle column shows sine minus periodogram scaled by periodogram error, and the right one shows sine minus Slepian scaled by Slepian error.

Fig. 12.— Four multitaper power spectra of $\ell = 65$ and $m = 0$ of the GONG month 16 velocity

time series (28 Oct 1996 – 2 Dec 1996, 36 days) for different number of tapers ($K = 5, 10, 20$, and 50).

Table 1: Number of modes fitted in three power spectrum estimates of the GONG month 16 velocity time series of $\ell = 30$, 65, and 100. The table shows the total number of modes and the number of good fits. The numbers of good fits are separated into three frequency ranges: (1) $\nu \leq 2.5$ mHz, low S/N modes; (2) 2.5 mHz $< \nu < 3.5$ mHz, well-resolved modes; (3) $\nu \geq 3.5$ mHz, blended modes.

$\ell = 30$	total	good			
	all ν	all ν	(1)	(2)	
periodogram	1196	516	143	187	186
Slepian	1146	914	277	314	323
Sine	1153	915	274	316	325
$\ell = 65$	total	good			
	all ν	all ν	(1)	(2)	
periodogram	2333	1389	215	598	576
Slepian	2260	1880	447	669	764
Sine	2263	1851	436	669	746
$\ell = 100$	total	good			
	all ν	all ν	(1)	(2)	
periodogram	3077	837	169	387	281
Slepian	3092	1658	545	535	578
Sine	3108	1628	509	544	575

Table 2: As Table 1, but for wavelet denoised spectra.

$\ell = 30$	total all ν	good			
		all ν	(1)	(2)	
periodogram	783	486	124	246	116
Slepian	1198	956	286	328	342
Sine	1199	959	292	332	335

$\ell = 65$	total all ν	good			
		all ν	(1)	(2)	
periodogram	1702	905	125	422	358
Slepian	2231	1677	451	663	563
Sine	2241	1681	457	666	558

$\ell = 100$	total all ν	good			
		all ν	(1)	(2)	
periodogram	2498	344	33	83	228
Slepian	3086	1561	567	552	442
Sine	3093	1554	575	550	429

Table 3: Linear regression of mode parameters (ν : frequency, Γ : width, and A : amplitude) between any two power spectrum estimates (cf. Figure 9) for $N = 1285$ good fits common to all three power spectra of $\ell = 65$ determined for all frequencies and separately for three frequency ranges: (1) $\nu \leq 2.5$ mHz, low S/N modes ($N = 203$); (2) 2.5 mHz $< \nu < 3.5$ mHz, well-resolved modes ($N = 593$); (3) $\nu \geq 3.5$ mHz, blended modes ($N = 489$). The table shows slope (a) and intercept (b). The intercept of frequency and width are in nHz.

$\ell = 65$	slepian /periodogram		sine taper /periodogram		sine /slepian	
	a	b	a	b	a	b
ν	0.909 ± 0.122	-9 ± 11	0.905 ± 0.136	-9 ± 11	0.981 ± 0.046	1 ± 11
(1)	0.982 ± 0.090	-8 ± 21	0.979 ± 0.100	-6 ± 21	1.011 ± 0.090	3 ± 22
(2)	1.006 ± 0.055	-5 ± 14	1.007 ± 0.068	-6 ± 14	1.011 ± 0.048	-5 ± 14
(3)	0.847 ± 0.178	-25 ± 45	0.840 ± 0.187	-32 ± 46	0.961 ± 0.069	-11 ± 43
Γ	0.862 ± 0.064	304 ± 128	0.725 ± 0.106	530 ± 221	0.989 ± 0.022	10 ± 54
(1)	0.802 ± 0.077	378 ± 97	0.787 ± 0.080	388 ± 100	0.989 ± 0.100	9 ± 139
(2)	0.935 ± 0.044	130 ± 82	0.930 ± 0.050	125 ± 93	1.002 ± 0.044	-18 ± 82
(3)	0.762 ± 0.138	1486 ± 1090	0.558 ± 0.197	2724 ± 1767	0.975 ± 0.055	159 ± 473
A	1.003 ± 0.018	-0.3 ± 2.4	0.987 ± 0.020	0.8 ± 2.5	0.983 ± 0.015	1.0 ± 2.2
(1)	0.954 ± 0.074	-3.3 ± 13.7	0.942 ± 0.075	-3.4 ± 13.7	0.989 ± 0.069	-0.2 ± 11.5
(2)	1.011 ± 0.040	4.1 ± 95.1	1.002 ± 0.041	-4.6 ± 97.7	0.994 ± 0.038	-14.1 ± 93.0
(3)	1.002 ± 0.023	0.0 ± 2.6	0.985 ± 0.025	1.3 ± 2.8	0.979 ± 0.019	1.4 ± 2.4

Table 4: Linear regression of mode parameters (ν : frequency, Γ : width, and A : amplitude) between the periodogram and the two denoised multitaper spectra (cf. Figure 10) for $N = 1226$ good fits (bad=0 and ierr=0) common to all three power spectra of $\ell = 65$ determined for all frequencies and separately for three frequency ranges: (1) $\nu \leq 2.5\text{mHz}$, low S/N modes ($N = 202$); (2) $2.5 \text{ mHz} < \nu < 3.5\text{mHz}$, well-resolved modes ($N = 587$); (3) $\nu \geq 3.5\text{mHz}$, blended modes ($N = 437$). The table shows slope (a) and intercept (b). The intercept of frequency and width are in nHz.

$\ell = 65$	slepian /periodogram		sine taper /periodogram		sine /slepian	
	a	b	a	b	a	b
ν	0.960 ± 0.097	-12 ± 11	0.957 ± 0.110	-9 ± 11	0.974 ± 0.052	2 ± 11
(1)	0.989 ± 0.093	-17 ± 21	0.981 ± 0.102	-10 ± 21	0.999 ± 0.089	7 ± 22
(2)	1.007 ± 0.060	-8 ± 14	1.010 ± 0.070	-7 ± 14	1.009 ± 0.049	1 ± 14
(3)	0.924 ± 0.146	-24 ± 45	0.918 ± 0.155	-14 ± 45	0.948 ± 0.083	5 ± 45
Γ	0.952 ± 0.063	170 ± 121	0.953 ± 0.071	144 ± 137	0.989 ± 0.022	-4 ± 54
(1)	0.828 ± 0.080	368 ± 100	0.804 ± 0.081	372 ± 102	0.975 ± 0.097	11 ± 136
(2)	0.945 ± 0.046	137 ± 86	0.936 ± 0.050	124 ± 94	0.995 ± 0.044	-20 ± 83
(3)	0.929 ± 0.148	523 ± 1102	0.930 ± 0.169	506 ± 1262	0.973 ± 0.052	179 ± 418
A	0.956 ± 0.019	-1.5 ± 3.6	0.959 ± 0.020	-1.0 ± 3.6	1.003 ± 0.016	0.3 ± 3.2
(1)	0.898 ± 0.071	-2.9 ± 13.4	0.912 ± 0.073	-3.1 ± 13.7	1.016 ± 0.071	-0.2 ± 11.6
(2)	0.965 ± 0.039	12.5 ± 91.3	0.976 ± 0.040	-1.4 ± 95.5	1.013 ± 0.038	-16.9 ± 91.2
(3)	0.952 ± 0.025	-0.4 ± 4.0	0.952 ± 0.026	-0.3 ± 3.9	0.998 ± 0.022	0.7 ± 3.6

Fig. 1

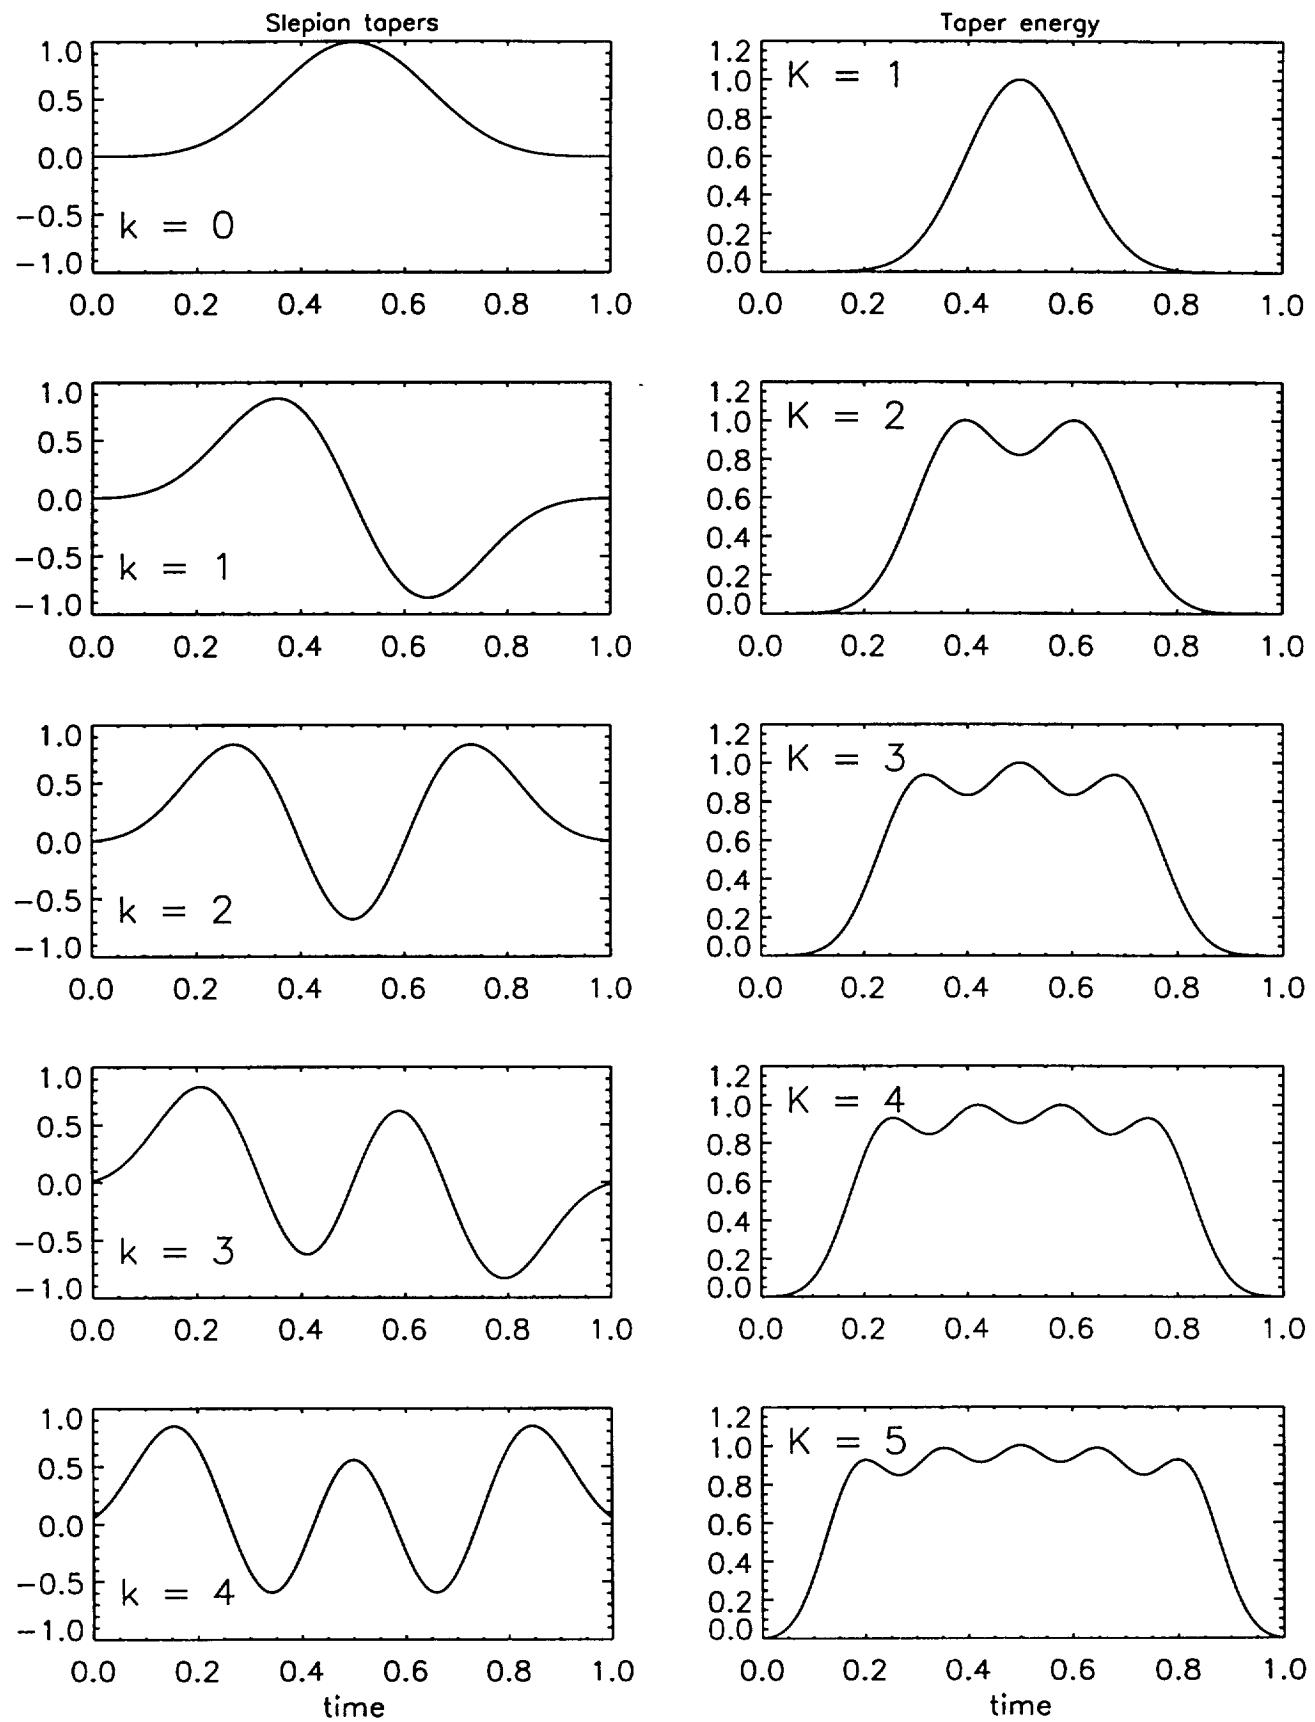


Fig. 2

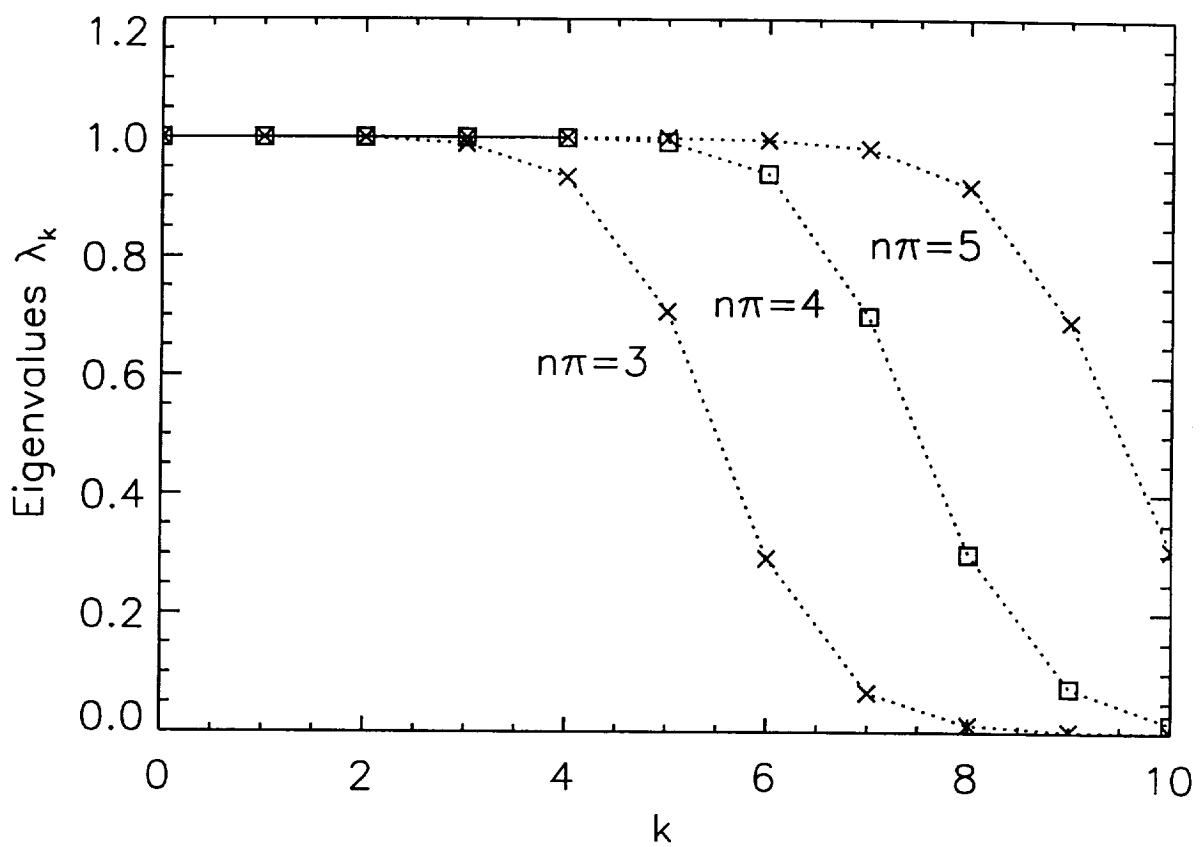


Fig. 3

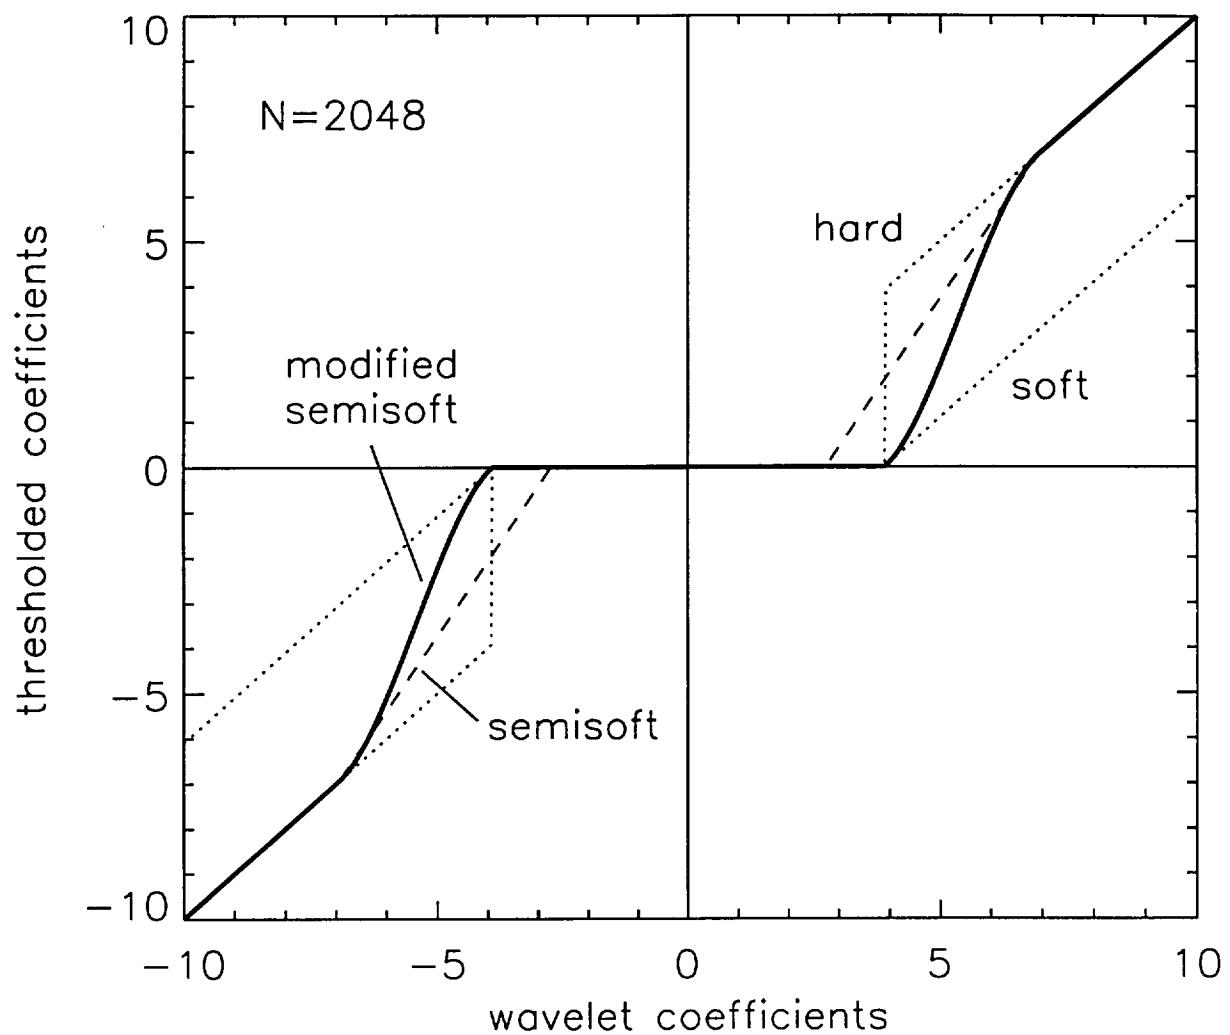


Fig. 4

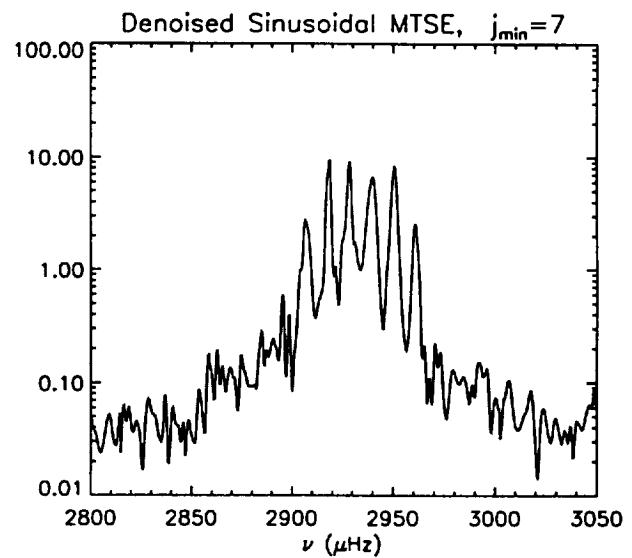
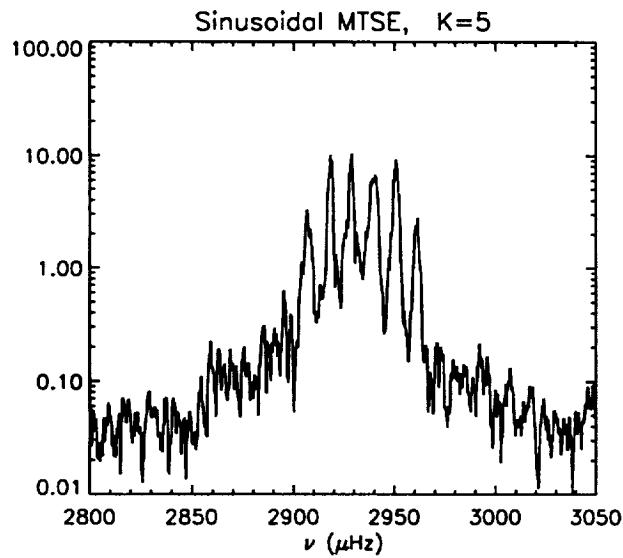
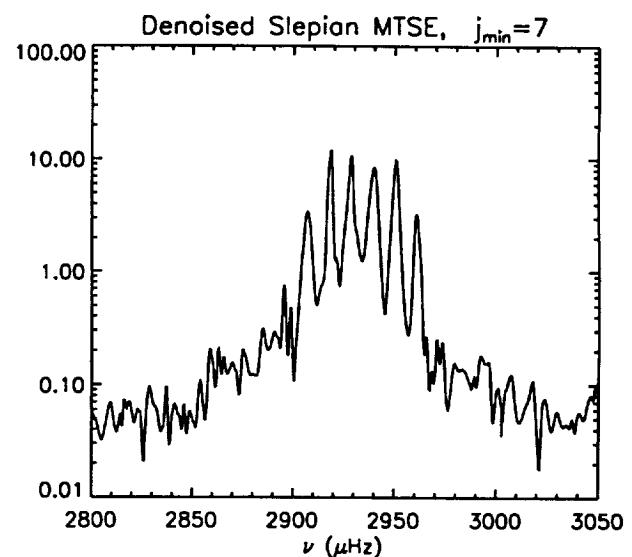
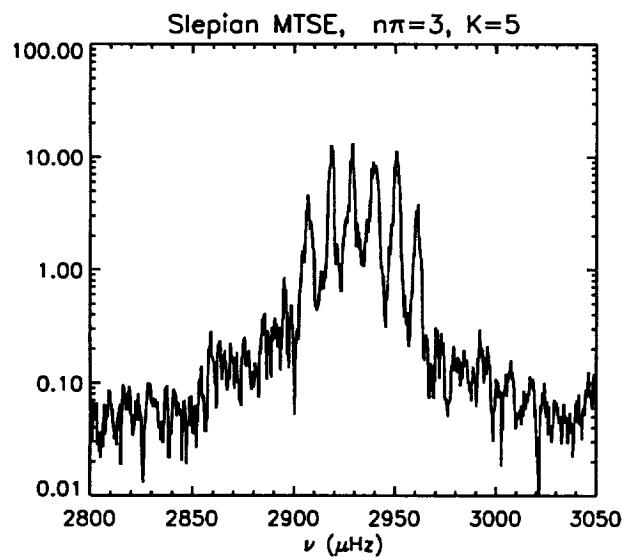
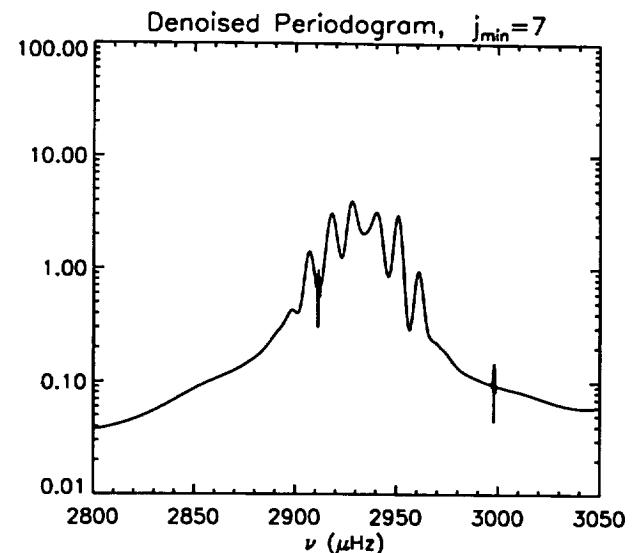
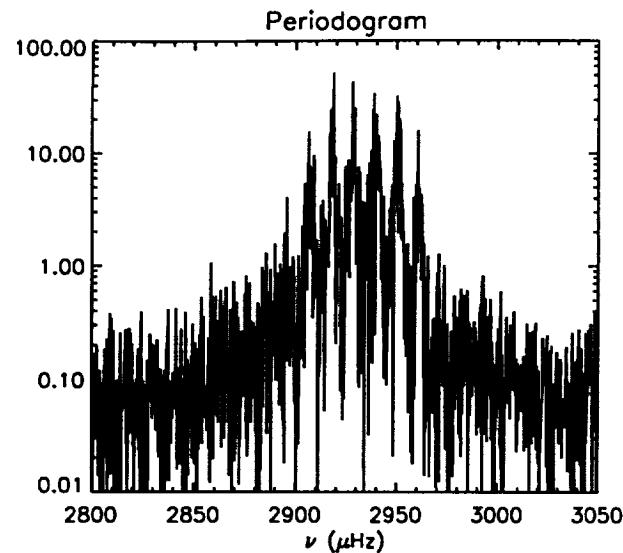


Fig. 5

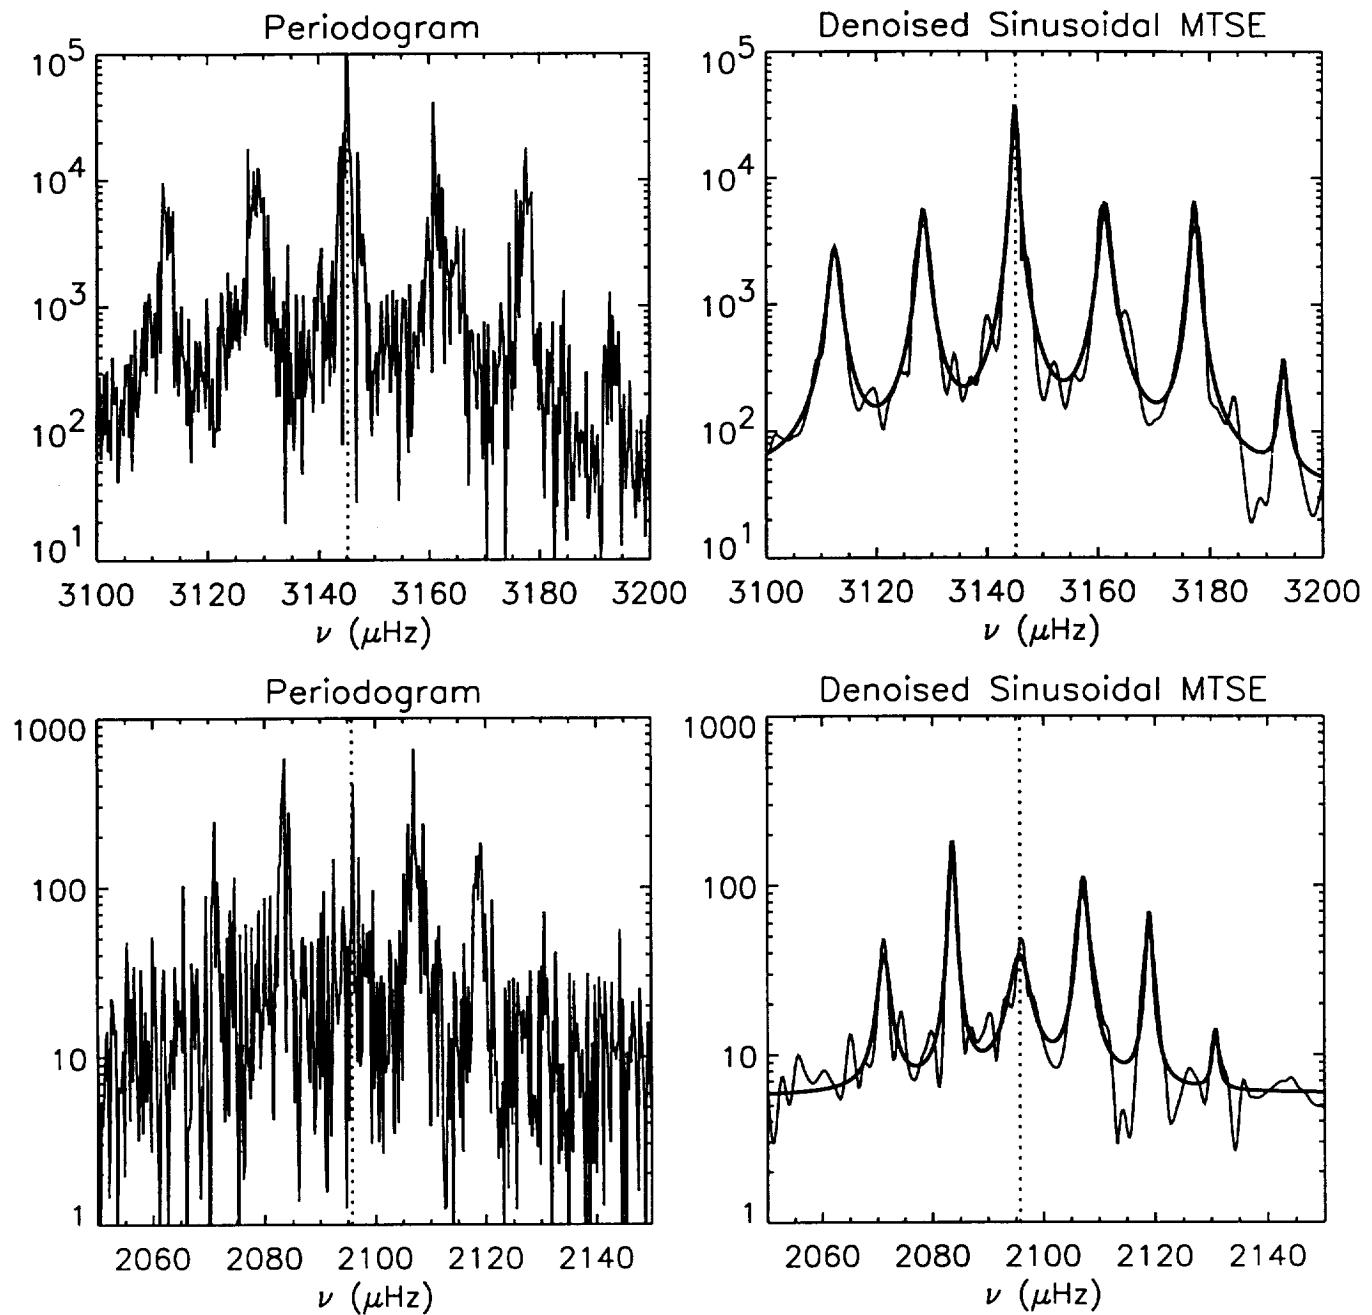


Fig. 6

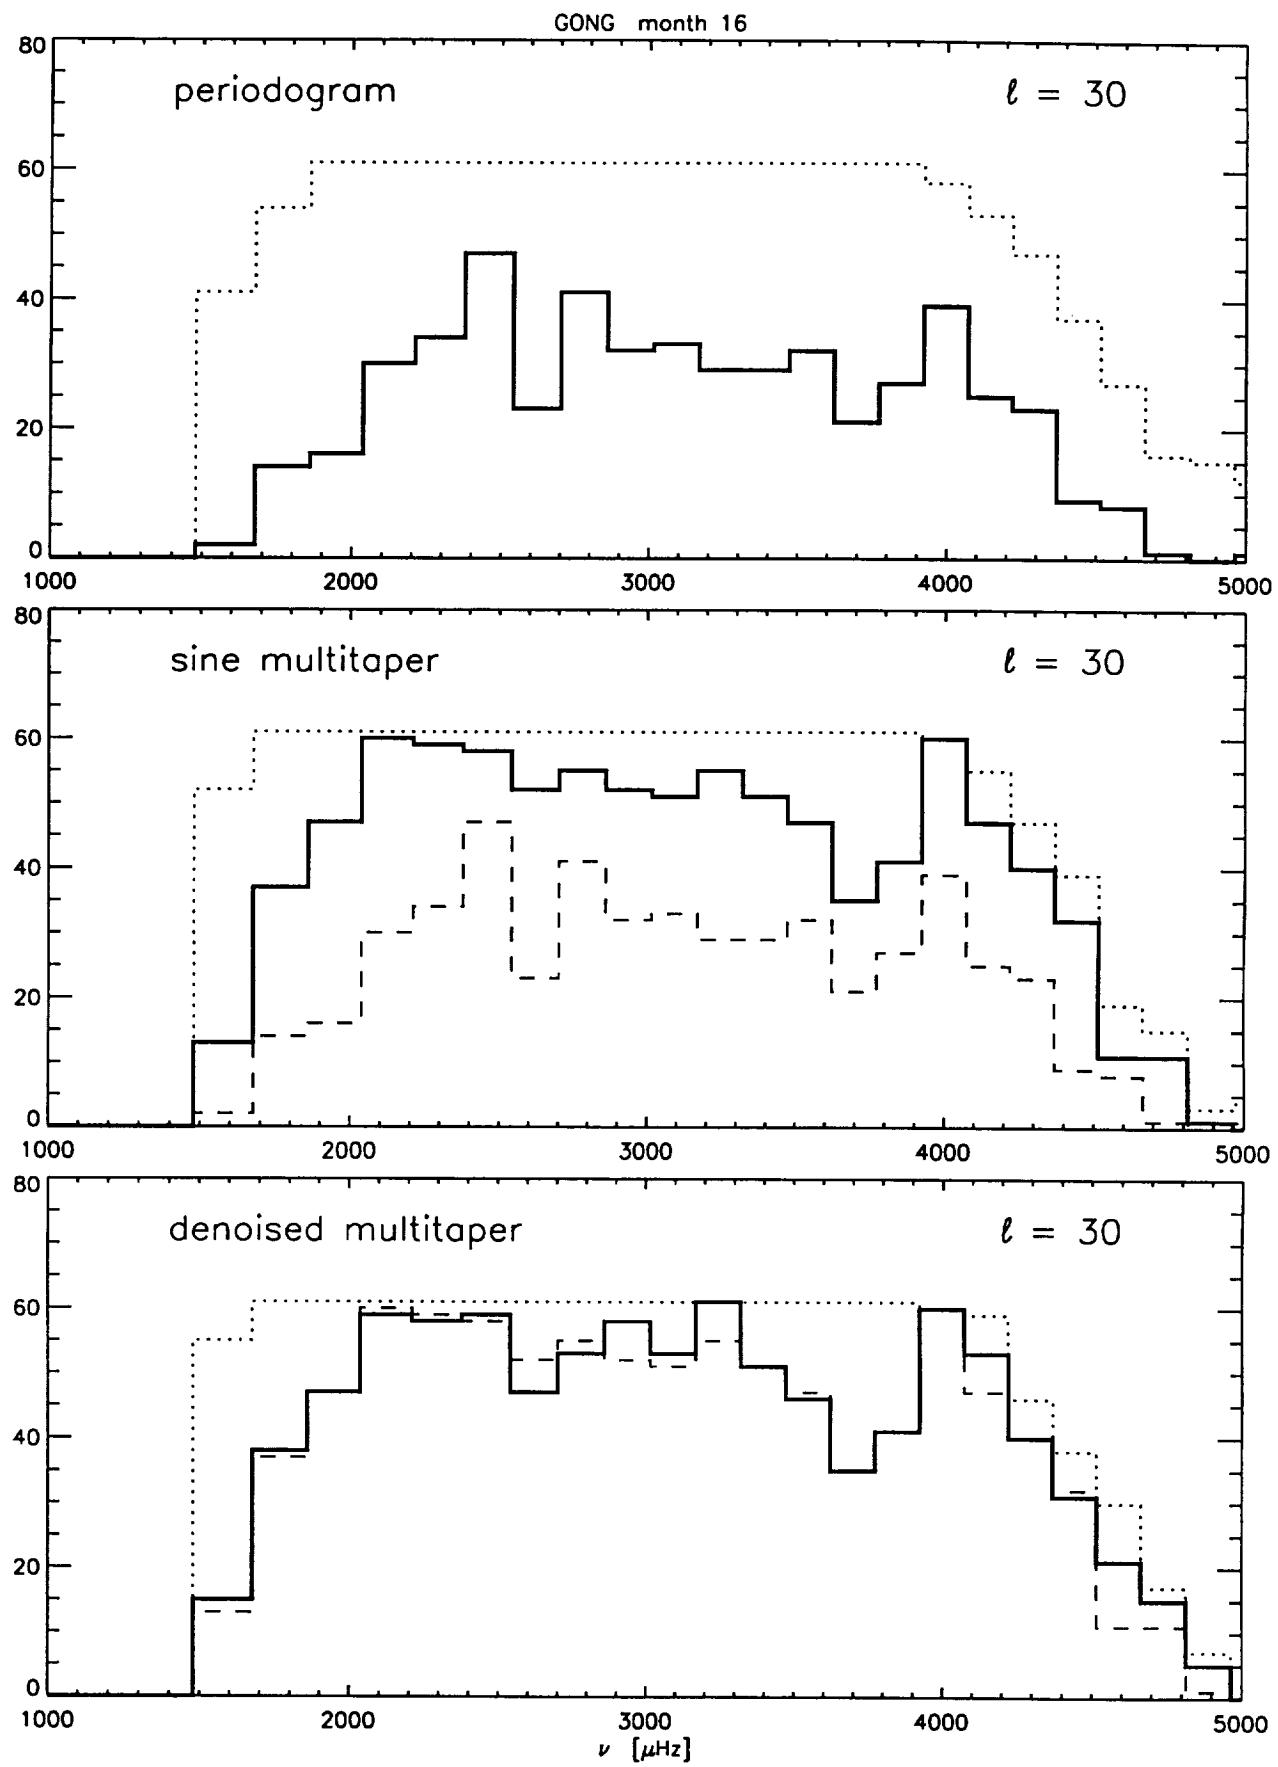


Fig. 7

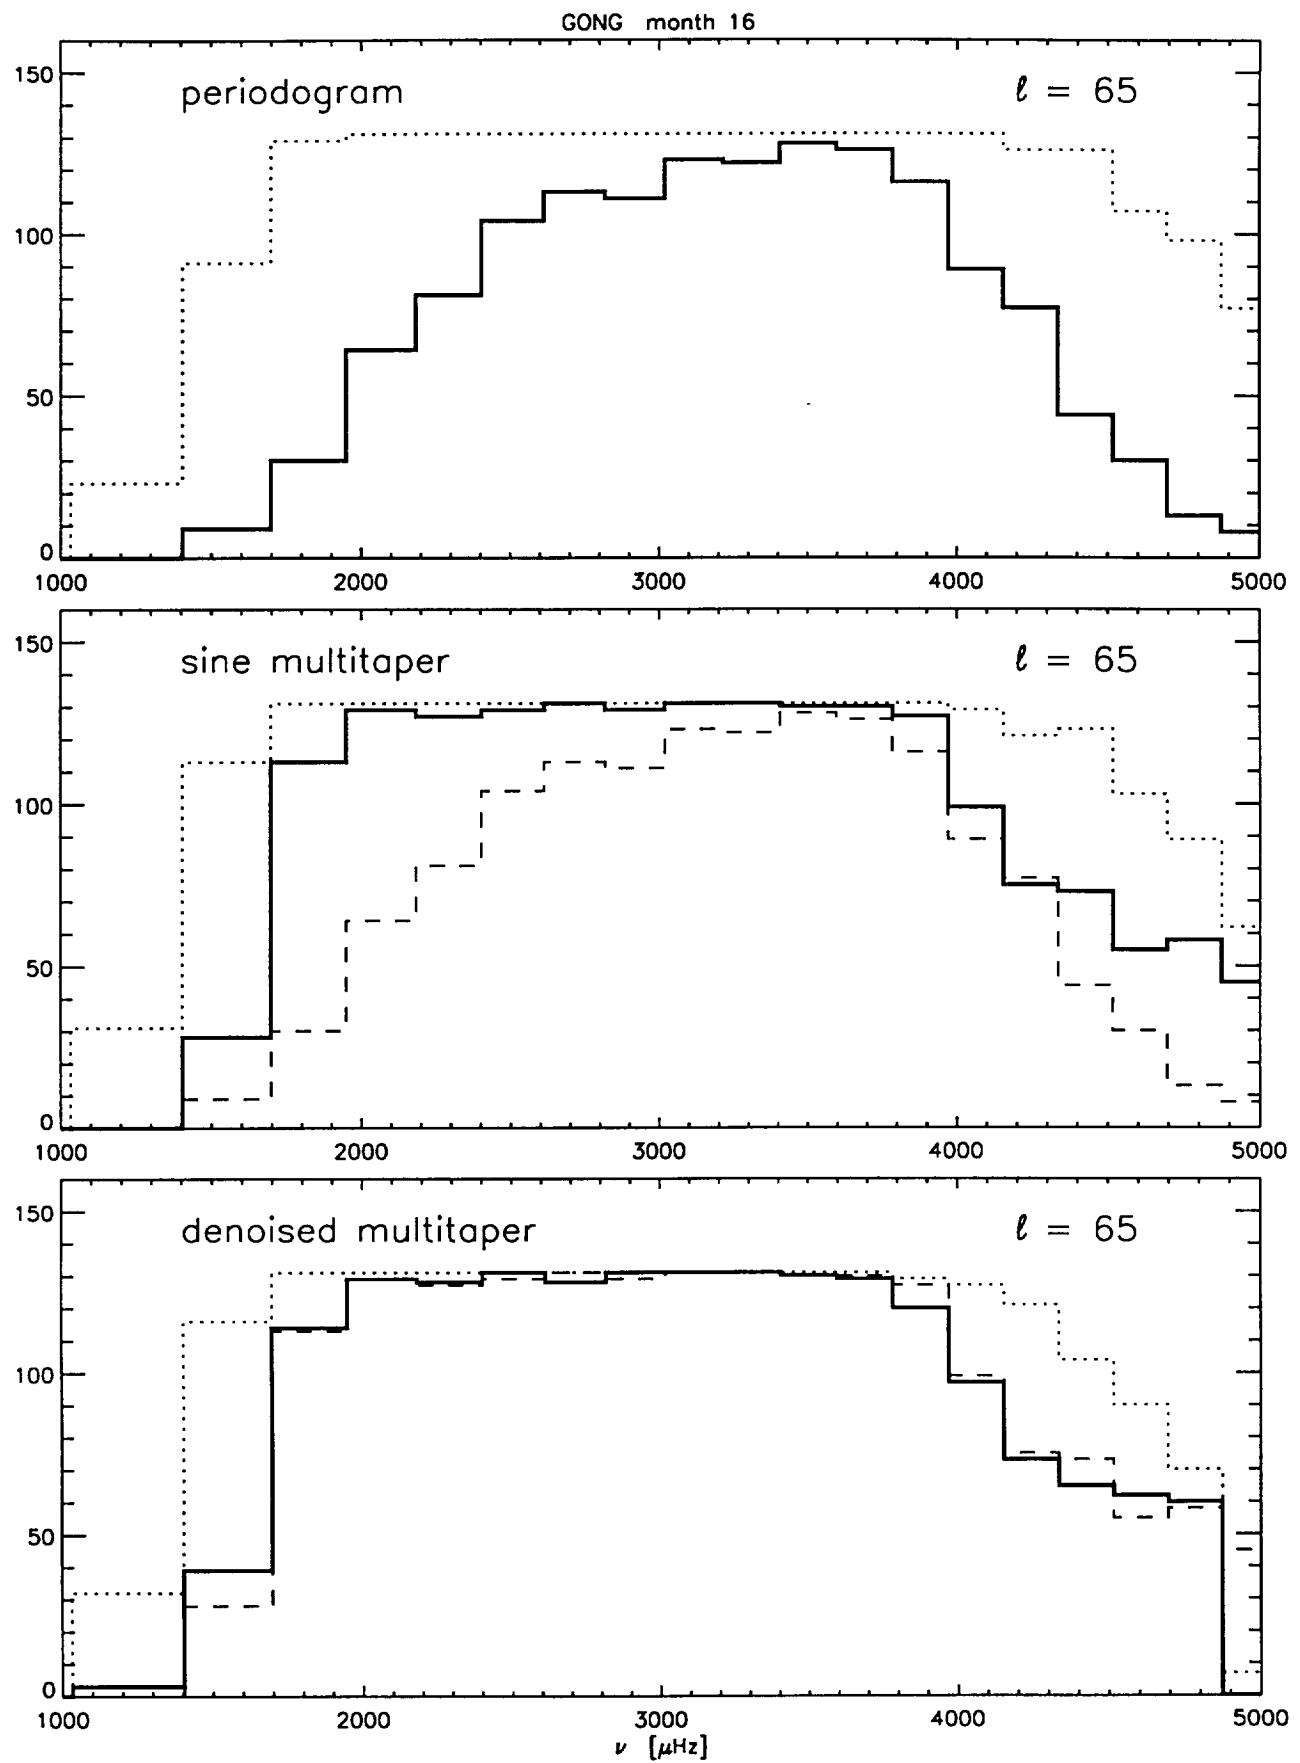


Fig. 8

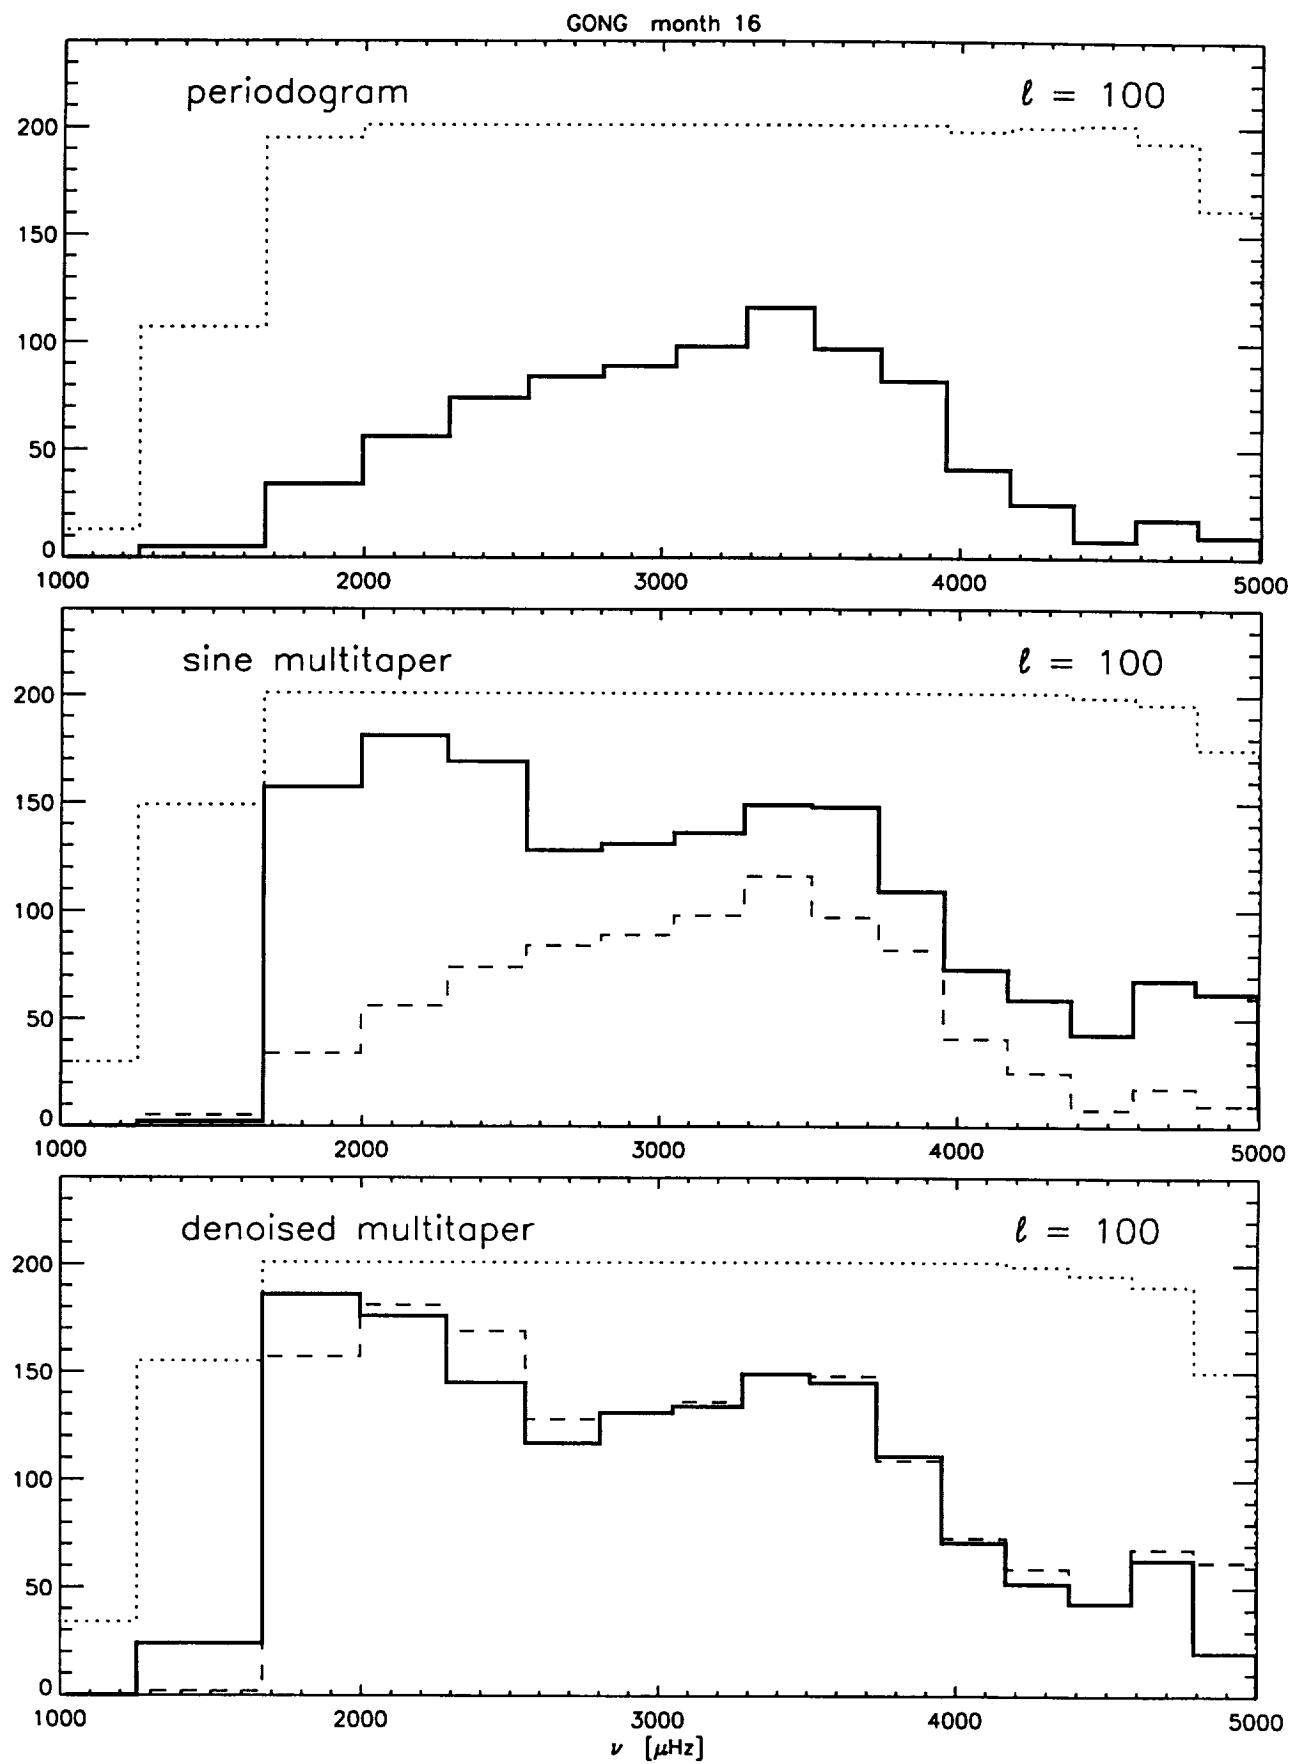


Fig. 9

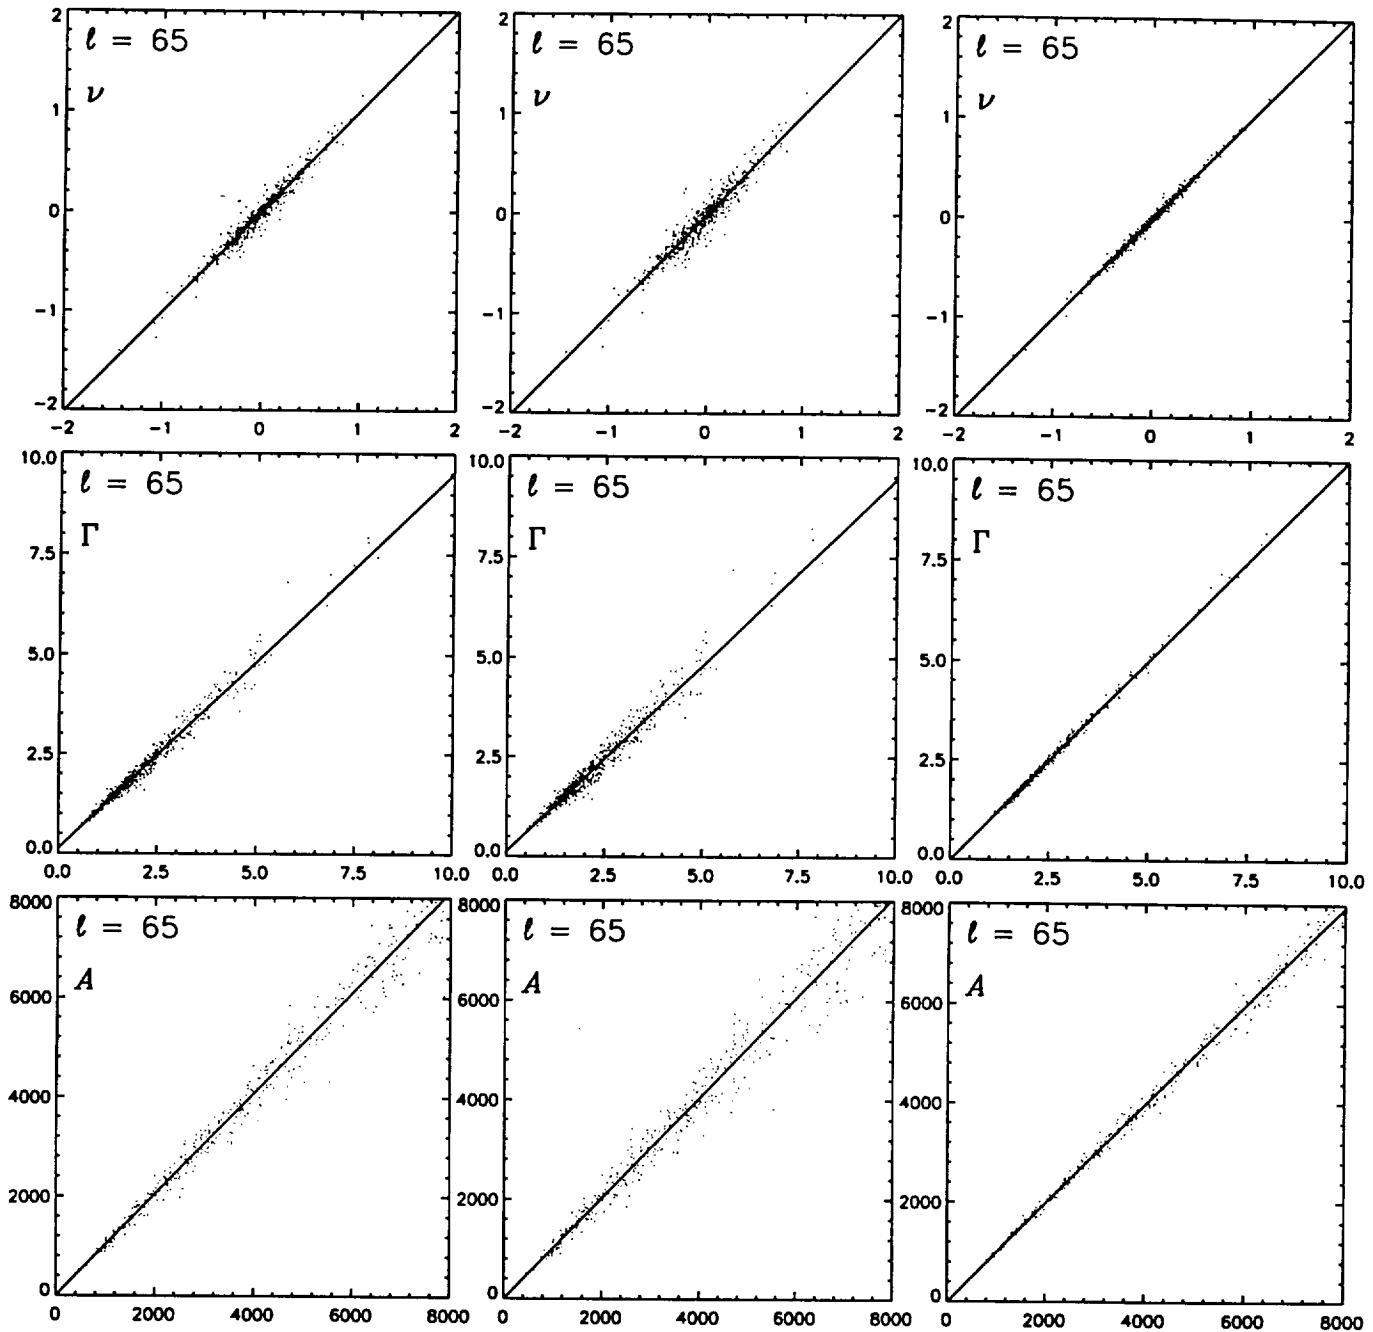


Fig. 10

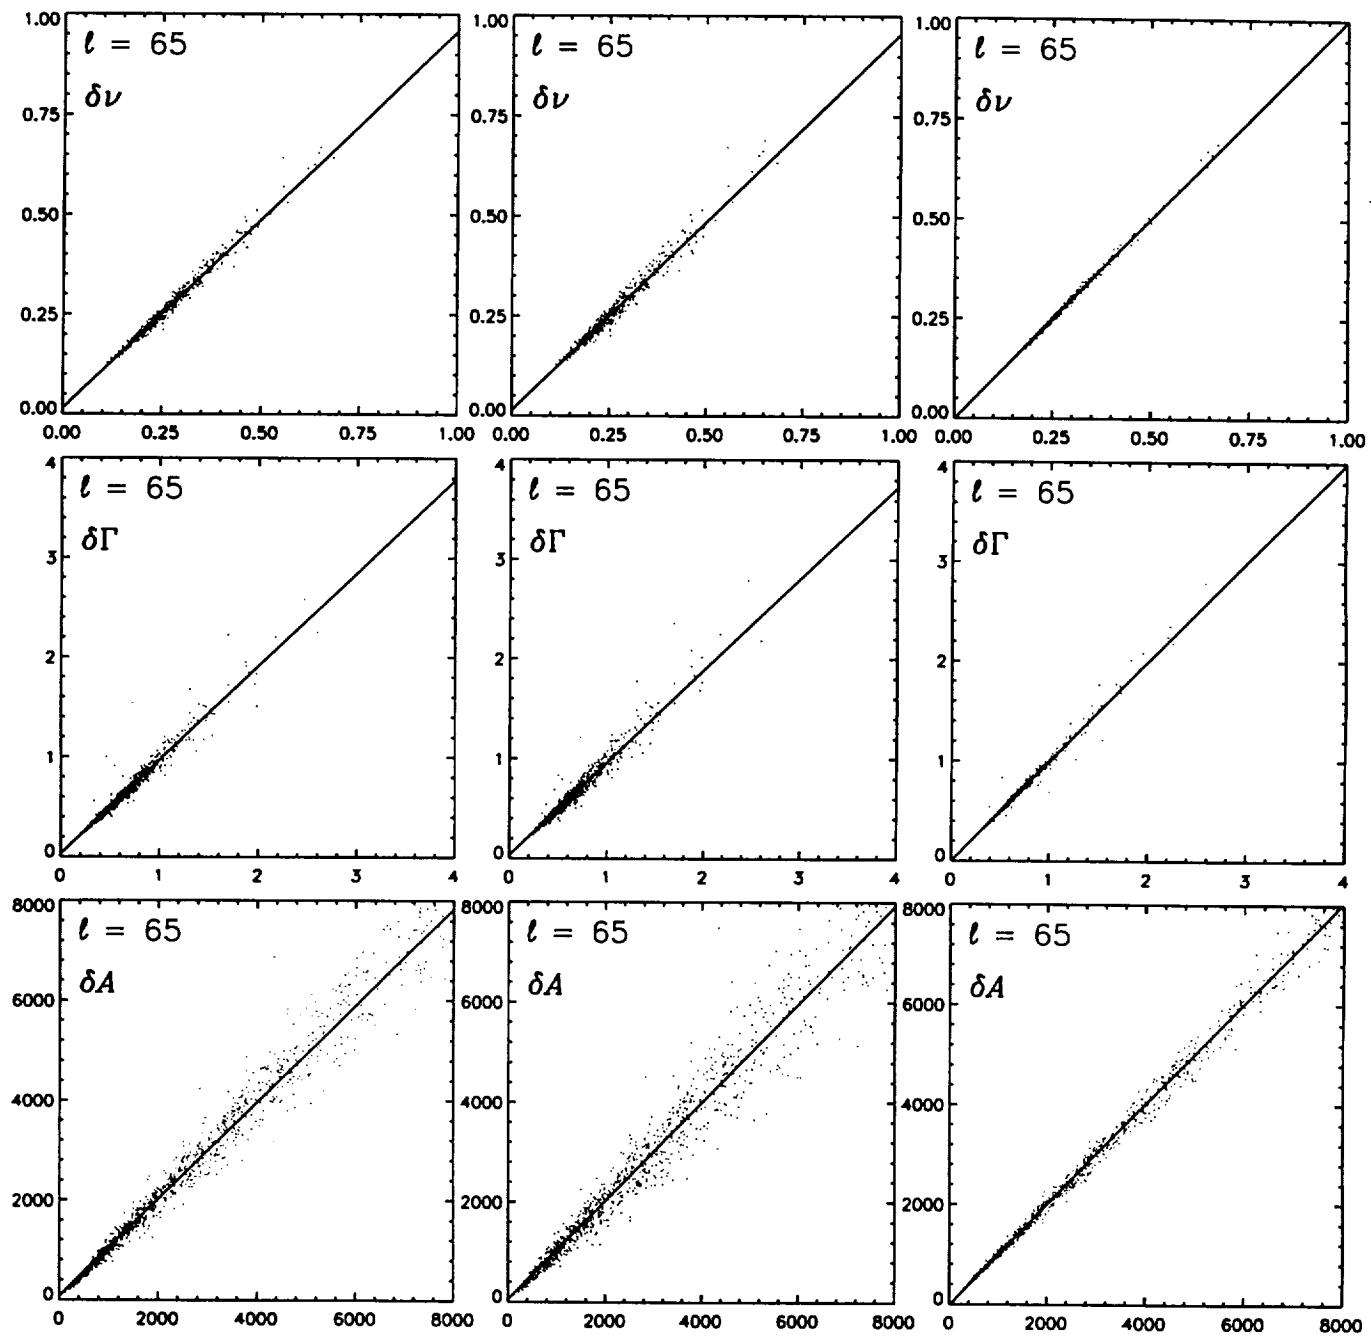


Fig. 11

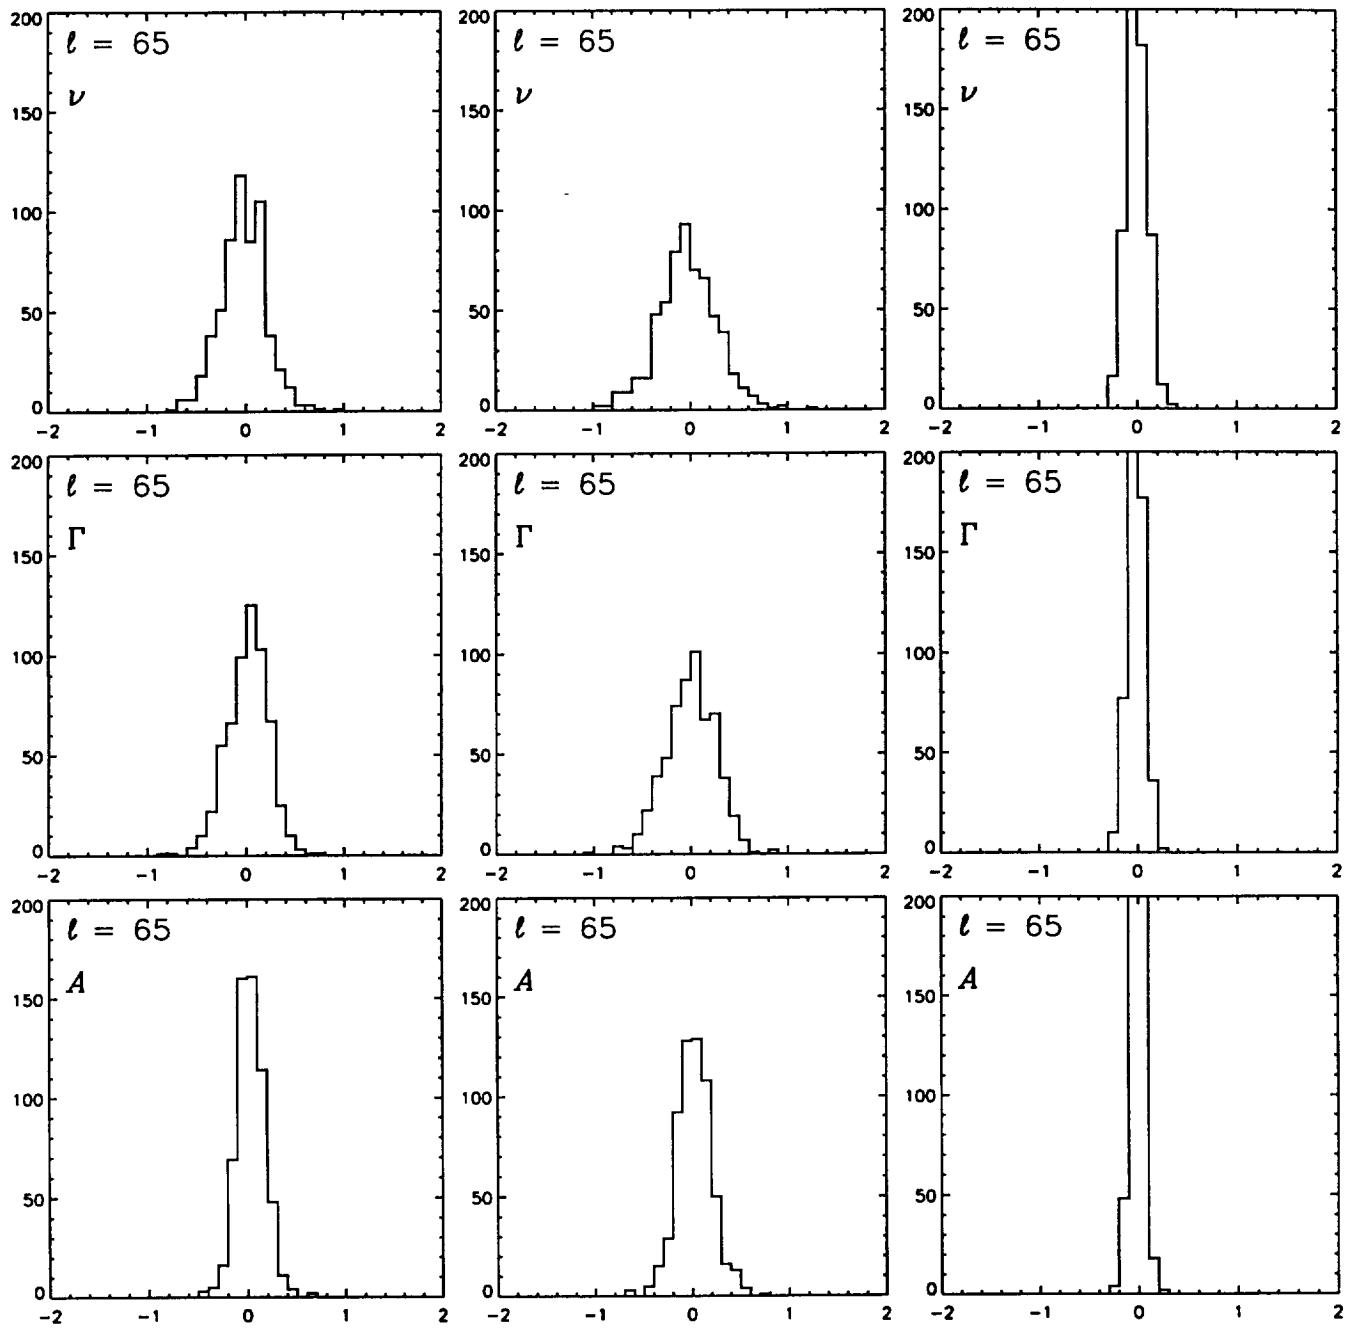


Fig. 12

